

Negative databases

The intriguing idea of a “negative database” that stores everything *except* the information of interest has been proposed recently (Esponda 2005) as a means for providing data privacy against *partially-specified queries*, while at the same time providing an efficient mechanism to answer *fully-specified queries*. To make this concrete, consider a database containing information about the names, institutions and professions of individuals. Here, partially-specified queries would look like “List all engineers in the database”, whereas fully-specified queries would be of the form “Does {P Balam, Indian Institute of Science, Molecular Biologist} exist in the database?”. In a nutshell, the goal of a negative database is to prevent “probing attacks” without adversely impacting standard verification queries. The crux of the technique is that although logically trivial, it is *computationally* extremely hard to invert a well-designed negative database and thereby determine the original contents.

A closed-world model is assumed wherein the universe U of all feasible entries in the database is known in advance, and the current database DB has a (typically small) subset of these entries. Today’s database management systems would store just DB , but in the proposed technique, what would be retained is the *complement*, that is, $\{U - DB\}$, resulting in the “negative database” name. Since $\{U - DB\}$ will usually be much, much larger than DB , an obvious difficulty is whether the negative database can be stored efficiently. This problem is solved through a well-known computer science representation called “regular expressions” wherein symbols such as “*”, “?”, “+”, etc. are used to denote a rich variety of constraints. For example, the regular expression “*ala*” is used to refer to all data entries that have the substring “ala” present *somewhere* in their value. So, if this regular expression was present in the names column of the negative database, it would automatically rule out all individuals whose names feature this substring – for example, Balam and Balakrishnan.

The second difficulty is whether it is possible *deterministically* to create a negative database that is robust to probing attacks. While this is computationally difficult in general, approximate techniques that provide such databases at the expense of incurring a few “false positives” – that is, a few additional data records not present in the original database will be incorrectly presumed to be present – have been proposed in (Esponda *et al* 2006).

From a computer science perspective, the negative database approach to supporting privacy in databases is novel and opens up challenging research questions. Interestingly, the inspiration for this approach is biologically rooted, as explained in the article “The non-denial of the non-self” (*Economist* August 31, 2006), where a parallel is drawn between negative databases and the human immune system. The analogy is that the immune system protects its owner from pathogens without knowing what a pathogen looks like, and the article goes on to say that “the immune system learns early on which biological molecules are “self”, and when it meets one that is “not self”, it destroys it”.

While the computer science aspects of negative databases are unexceptionable, there appears scope for debate about the extent to which the biological analogy is a meaningful comparison. On the one hand, it certainly appears reasonable to view the collection of immune cells as a negative database since they are derived by eliminating those that bind to self. On the other hand, there are a variety of problematic issues: For example, a question that arises immediately is whether the closed-world model of the negative database, which assumes that all feasible entities are known in advance, is practicable in the immunological world, where there is much uncharted territory. However, this is not a particularly troubling issue since the negative database can always be extended to accommodate new biomolecules as and when they are identified.

What is of more relevance is the immune system model itself – in recent times, the very basis of the classical “self vs non-self” model is being questioned. Some biologists now hold that the immune system

Keywords. Database; immune system

often uses explicit markers of non-self, called Pathogen Associated Molecular Patterns, to recognize non-self antigens (Janeway 1989). Even more radically, there are also proponents of the theory that "...the immune system is more concerned with damage than foreignness" (Matzinger 2002), i.e. it is injury that triggers immune reactions, not whether the foreign organism passes a shibboleth test.

Finally, a negative database is meant for keeping out nosy outsiders – in the biological world, the analogy would be a malicious protein that wants to check out what types of molecules are considered "self" by the immune system. It is not clear how such proteins would be able to either pose a sequence of "partially-specified" queries to the immune system in order to break down its defences, or non-invasively obtain the information encoded in the negative database of the immune cells, especially given that each question could potentially result in the death of the questioner.

In closing, the negative database concept simultaneously opens up challenging research problems and piquant connections with natural biological processes.

References

- Esponda F 2005 *Negative representations of information*, PhD Thesis, University of New Mexico, Albuquerque, New Mexico, USA
- Esponda F, Ackley E, Helman P, Jia H and Forrest S 2006 Protecting data privacy through hard-to-reverse negative databases; *Proc. of 9th Information Security Conf.*, September 2006
- Janeway C A Jr 1989 Approaching the asymptote – Evolution and revolution in immunology; *Cold Spring Harbor symp. Quant. Biol.* **54** 1–13
- Matzinger P 2002 The Danger Model: A renewed sense of self; *Science* **296** 301–305

JAYANT R HARITSA
Supercomputer Education and Research Centre
Indian Institute of Science,
Bangalore 560 012, India
(Email, haritsa@dsl.serc.iisc.ernet.in)

ePublication: 13 November 2006