

A Network Use for Incomplete Knowledge Management

Anduo Wang (anduo.wang@gmail.com), Temple University, USA

Data management as a means to enable innovation in the networking domain was once highlighted in declarative networking [4], but the network use for databases in general has been a mixed success. Onix, a closed-source software defined network (SDN) platform for large-scale production networks that marked a milestone in the evolution of SDN, achieves strong consistency among network states by replicated transactional databases. These heavy-weight databases, however, were abandoned in Onix’s decedent (the open-source ONOS), and gave way to a distributed key-value store that minimizes index maintenance and read-write operations. Similarly, in the high-profile network verifier Batfish, the datalog engine used to track data provenance in the original version was later replaced by an imperative one for scalability. In these unfortunate recurring developments, what the distributed system and imperative programming lacked at first in elegance (e.g., declarative abstraction, strong guarantee), they make up in performance and scalability in the long run.

Can database make a more permanent impact in networking? We argue for the affirmative and to achieve so we suggest a new use for database, the management of incomplete network knowledge. Incomplete knowledge has been inherent, prevalent, and will likely to persist in computer networks; it is also a problem that cannot be (easily) addressed through means other than database.

At the *data level*, incomplete knowledge naturally arises as what network admins really care about is not the behavior of a specific network, but to ensure an acceptable operating region — such as all possible network behaviors permitted by a topology meeting a certain criteria (e.g., not leaking traffic). Existing tools (e.g., formal network verification that gains significant attention recently) focus on network state (e.g., forwarding configuration distributed at individual routers that collectively determine packet delivery) while at a particular instance. Such mismatch — between the incomplete operating region requirement and the complete instantaneous analysis — cannot be overcome by improving existing approaches (e.g., faster verification algorithm or optimized representation of complete network instance). What is needed is genuine support for incomplete network states, treating them as definite data. To our best knowledge, only in incomplete database [2] (e.g., the *c*-tables) is incomplete data managed as first class citizen.

At the *semantic level*, network admins have been struggling to understand the properties a network design targets, to predicate what the implementation actually produces, with limited visibility into the network — in the public Internet, the ownership and administration of, and thus visibility into, the Internet has been divided into individual and oftentimes competing domains. This leaves any attempt to verify constraints on the global Internet that requires a complete view futile. Indeed, constraint checking in the Internet — even for basic property such as connectivity under link

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution, provided that you attribute the original work to the authors and CIDR 2022. 12th Annual Conference on Innovative Data Systems Research (CIDR’22), January 9-12, 2022, Chaminade, USA.

failures — remains a black art. This is in contrast to distributed databases where partial information has been explicitly accounted for during integrity constraints maintenance.

At the *structure level*, the ever evolving network protocol stack lacks a well-defined control flow. After decades of unplanned protocol deployments that show no sign of stopping, a computer network is better described as numerous distributed facts that take heterogeneous forms, that engage in intertwined dependencies, and that interact in complex and unforeseeable ways. Existing efforts to bring structure to network data so as to simplify its management remains isolated and fragmented. Since the complexity of networks is unparalleled in software development, the challenge of finding a network schema — one that fits all needs and accommodates requirements of the present as well as the future — will likely outmatch what the software world can offer. In fact, rarely do we see a more elegant solution to irregular data than what the semistructured data research [1] already offered.

	DB solution	networking use
data	incomplete database	operating region management
semantics	distributed database	inter-domain routing
schemas	semistructured data	network coverage

Table 1: A network use for database: incomplete knowledge management at the data-, semantic-, and schema- level.

This abstract calls for the database community to make a more permanent impact in networking, and presents a first of its kind network use for incomplete knowledge management. Incomplete knowledge which has been a pain point for non-data approaches, is particularly relevant for networks, at every level of data management, as summarized in Table 1: Network admins care more about the partially specified operating region instead of a completely known instantaneous state. Our recent work [3] demonstrates a promising solution with incomplete database in which the *c*-table gives a natural representation for operating regions, the datalog semantics can be cleanly extended to account for precise operation, and the PostgreSQL allows practical implementation. Beyond incomplete data, we illustrate how the long-standing verification problem of the Internet maps to constraint checking in distributed databases, and how the emerging network coverage problem can be formulated and significantly enhanced by semi-structured data querying.

Acknowledgments. This work was supported by National Science Foundation Award CNS-1909450.

REFERENCES

- [1] S. Abiteboul, D. Quass, J. McHugh, J. Widom, and J. L. Wiener. The lorel query language for semistructured data. *International Journal on Digital Libraries*, 1997.
- [2] T. Imieliński and W. Lipski. Incomplete information in relational databases. *J. ACM*, 31(4):761–791, Sept. 1984.
- [3] F. Lan, B. Gui, and A. Wang. Fauré: a partial approach to network analysis. HotNets’21.
- [4] B. T. Loo, J. M. Hellerstein, I. Stoica, and R. Ramakrishnan. Declarative routing: Extensible routing with declarative queries. SIGCOMM ’05.