

Letter from the Editor-in-Chief

How to efficiently and effectively manage large-scale data is a critical challenge in data management, scientific computing, machine learning, and many other fields. In this issue, we look into this problem from two angles.

Gerhard Weikum’s opinion piece titled “Entities with Quantities” highlights development along the direction of querying the Web as a database. We have come a long way in keyword based Web search: Today, all major search engines support entity based question/answering to certain extent (e.g., returning “Eiffel Tower” for query “the highest building in Paris”). Weikum is taking one important step towards the goal of querying the Web as a database. In the article, he discusses what it takes to find all entities that satisfy a quantity-based search condition, for example, “buildings taller than 500m” or “runners completing a marathon under 2:10h.” It is clear that this requires much advanced data preprocessing (e.g., information extraction, entity linking, etc.), but more importantly, it requires that at least part of the data on the entire Web needs to be organized as a database.

Philippe Bonnet put together the current issue consisting of 5 papers from leading researchers in the high performance computing and data management communities on the topic of data management at Exascale. Advances in exascale computing on petascale supercomputers are pushing the frontier of scientific computing that requires complex simulation, benefiting applications ranging from astrophysical discovery to drug design. But with increasing amounts of data, the gap between computation and I/O has grown significantly wider, which makes data management a big challenge. This timely issue answers many questions in this domain.

Haixun Wang
WeWork Corporation