This issue of the Data Engineering Bulletin is devoted to the topic of data quality. Data quality is an important topic of vast business importance. Quality problems significantly degrade common business practices (with billing being a prominent example). Typical quality problems include incorrect person or business names, inconsistent address information, poor integrity constraints, missing keys to name a few. It is estimated that data quality problems cause approximately 600 billion dollars of financial loss in the US alone each year. Data quality has been in the center of research interest for more than thirty years. Although several steps have been made towards addressing major problems and an array of industrial tools are available, fundamental problems remain still widely unsolved.

The issue contains seven articles authored by leading experts in the area. The specific collection of articles was chosen in a way to highlight current research directions in data quality as well as new interesting approaches to known problems.

The first two articles present new approaches to the problem of *entity resolution*. The article by Bhattacharya and Getoor addresses the problem of collective entity resolution, namely exploring relationships or ties between entities in a way that related entities are resolved collectively. The article by Benjelloun et. al., is an overview of the approach to entity resolution by the Stanford SERF project. In this approach match predicates are considered as black boxes and algorithms are presented to merge entities efficiently.

The article by Naumann et. al., presents an interesting approach to data fusion by considering a multi-level approach. In particular the authors present the results of their research to merging data from different sources by systematically tracking inconsistencies at the value, tuple as well as schema level. Laure Berti-Equille discusses quality aware query processing. This is a relatively new and highly important problem. Traditional approaches to data quality consist of techniques to 'clean' data in order to enable query processing on the clean image of a data set. Quality aware query processing aims to enable query processing on 'dirty' data (data with various quality problems) and provide tradeoffs between query processing cost and result quality.

Dasu et. al., present work in the context of the Bellman data quality tool from AT&T Research. A very interesting direction in data quality deals with data base exploration for quality problems. It is highly important to provide tools and methodologies to identify data quality problems, in addition to providing techniques to resolve them once they have been identified. The Bellman tool is an important step in this area. The article explores data mining on database dynamics. In particular is performs data mining to explore database changes as a function of time. Chaudhuri et. al., present work in the context of the Data Debugger project at Microsoft Research. The main goal of the project is to identify generic and robust abstractions for data cleaning operators and to support efficient implementations of these abstractions. Doing so will enable composability of operators (relational and non relational) to derive operator trees. In the paper the authors illustrate their approach using the record matching operation.

Finally the article by Milano et. al., deals with addressing quality problems in non relational data sources. In particular it presents approaches for resolving identities for XML data taking into account the specifics of such data and in particular the underlying document structures.

I would like to take this opportunity to thank the contributors to this special issue once again for the time and effort they have put in delivering the articles. I would like to also thank the editor in chief of the publication Dr. David Lomet for his availability and timely response and handling of all matters related to the issue. Mr. Dimitris Tsirogiannis provided valuable assistance in assembling the issue and I would like to thank him for his time. I hope readers will enjoy the article collection as much as I have.