Foreword

This is the 9th workshop on this theme held annually. Traditionally, the workshop has been held along-side the SIAM datamining (SDM) conference, even if the first four editions were organized in conjunction with IPDPS. Over the years the definition of high performance computing has taken on various forms as a function of the types of technical and creative uses and the underlying semantics of the applications driving them. Traditional definitions often refer to the problem of using high end parallel computers to meet the need of scientific applications. However, high performance computing can also include the need for fast sequential algorithms that target memory and I/O performance. In this edition of the workshop, a wide variety of topics is covered. To begin, we are very honored to have Jiawei Han giving the keynote talk on "Warehousing and Mining Massive RFID Data Sets". Next, we were able to accept three contributed papers and three invited papers, representing a wide variety of topics, from efficient mining algorithms in high-speed data streams, distributed and grid-based environments to mining on emerging architectures and the acceleration of data mining workloads. We would like to thank Jiawei Han and all authors for their contributions, and the program committee (Gagan Agrawal, Ian Davidson, Sara Graves, Ruoming Jin, Hillol Kargupta, Shonali Krishnaswamy, Vipin Kumar, Salvatore Orlando, Raffaele Perego, Krishanmoorthy Sivakumar, Domenico Talia, Pang-Ning Tan and Mohammed Zaki) for helping us in the task of making this workshop again a success.

Bart Goethals
Byung-Hoon Park
Srinivasan Parthasarathy