One of the fundamental changes the Web will bring to us is the ability to organize, manage, discover, compose, and invoke *web services*. Speaking loosely, web services are bite-sized pieces of software system components that can be executed over a network (e.g., the Internet). An immediate and laudable goal is to be able to dynamically share web services in a similar manner to the sharing of data on the Web. To this end, the principles of Service Oriented Architecture (SOA) offer useful guidelines for constructing web services, compositions, and applications based on web services.

Research on web services started more than a half dozen years ago and, in recent years, has attracted increasingly more attention from several research communities including the database community. It frequently reminds me of the data management research prior to the arrival of the relational data model. Back then, there was not only a lack of a simple, yet commonly accepted data model, but also the prevailing models (network, hierarchical) were very much motivated by physical storage structures for data. Indeed, the search for a suitable common model for web services is still ongoing in spite of various standards (e.g., the ones with names matching "∗WS∗"). An important reason for this is insufficient understanding about web services, their description, use, composition, management, their interaction with stored data, etc. (SQL standardization happened years after significant research efforts were spent on the relational data model, query languages, query optimization, database design, etc.) Of course, history doesn't quite repeat itself. Models for web services emerge from the physical world (software development), as well as from many different research communities with different twists, adding abstractions, semantics, data, orchestration/choreography, etc., or their combinations.

Many research challenges are in front of us. Here are three examples. Automation, as a key aspect of SOA, requires incorporation of *semantics* in various phases of web service development and applications. Providing semantics in suitable form is the first challenge. The standard SAWSDL enables the capture of semantics about how the input and output arguments of a web service are linked to underlying domain ontologies. However, research on the semantics of the actual impact of services and how they are combined (e.g., OWL-S, WSMO, FLOWS) continues to be an area of active investigation. A machine-readable semantics that can be reasoned about can ensure *correctness* and at the same time help to achieve *efficient* services. Service composition is another area full of interesting research problems. There are many possible models for composition. In one extreme, we might write down a detailed *orchestration* of a collection of web services. In the other extreme, we could start from a global *choreography* that defines the expected behavior of the entire composition. Between the two extremes, one could figure out behavior interfaces of two or more services (e.g., as "embedded" or partial choreography), and gradually expand to an entire composition. The exact relationship among the composition models remains to be explored. As the third example of challenging areas, techniques for analyzing, verifying, and mining web services are in high demand. These will become especially important as the services used are increasingly created through semi- and fully-automated discovery and composition mechanisms.

Presented in this special issue is a sampler of ongoing research efforts in web services. By no means do they represent a complete survey of current projects in the area, but I hope they can be the starting point of your exploration in web services. As you will find, there are many practical motivations for this work, but no lack of technical depth and elegance. I recall that Tim Berners-Lee in his WWW '03 invited talk identified data and (web) services as two fundamental elements flowing on the Web, interacting and bearing "fruits", freely and effectively. In my view, most of the work to date treats the service/process aspect and the data aspect in largely separated ways; the fundamental interplay of data and services, and the development of new models that bring that interplay to the forefront, remain largely unexplored and untapped. Maybe some of you will find deep and engaging challenges for your skills and talents by contributing to this exciting (and highly profitable!) grand challenge in software development and service management.