Where are the Interesting Problems?*

Barry O'Sullivan

Cork Constraint Computation Centre
Department of Computer Science, University College Cork, Ireland
b.osullivan@cs.ucc.ie

Abstract. Constraint programming has become an important technology for solving hard combinatorial problems in a diverse range of application domains. It has its roots in artificial intelligence, mathematical programming, operations research, and programming languages. In this talk we will discuss a number of challenging application domains for constraint programming, and the technical challenges that these present to the research community.

1 Introduction

Constraint programming (CP) is a technology for solving combinatorial optimisation problems [3]. A major generic challenge that faces CP is *scalability* [2], largely because the problems to which it is usually applied are computationally intractable (NP-Hard). While CP has been successfully applied in domains such as scheduling, timetabling, planning, inventory management and configuration, many instances of these problems are extremely challenging for traditional CP methods due to their hardness.

However, an emerging dimension of scale relates to problem size, and the volume of data available that is relevant to solving a particular instance, e.g. extremely large domain sizes, or very large extensionally defined constraints of high arity. In 2009 information on the web was doubling every 18 months. It is now believed that this occurs in less than 12 months. This exponential growth in data, often referred to as the "big data" challenge, presents us with major opportunities. For example, McKinsey Global Institute estimates that European government administrations could benefit from over €250 billion in operational efficiencies by properly exploiting "big data".

Another major challenge in real-world application domains is *uncertainty* [4, 5]. For example, in scheduling, the duration of a task might be uncertain, while in inventory management there might be uncertainty related to customer demand. Surprising, even predicting the future does not imply that we can make better decisions. The interactions between the choices that face us are usually interlinked in complex ways. Being able to react appropriately to risk is more important than knowing about the risk or even modelling it. The traditional "get data – model – implement"–cycle is no longer sufficient in most domains.

* This work was supported by Science Foundation Ireland Grant 10/IN.1/3032.

We often need to deal with large amounts of rapidly changing data whereby adaptation becomes key. The study of managing complex sources of data upon which we must make complex, risky, economic or environmentally important, decisions provides a compelling context for constraint programming [1].

2 Challenge Domains

A number of important application domains are emerging due to the growth in large-scale and complex computing infrastructure, a desire to improve business efficiency, the need for improved quality in public services, and demands for sustainable environmental planning. In this talk we will discuss several examples including: data centre optimisation, e.g. energy management through workload consolidation; innovative enterprise and public service delivery, e.g. optimising access to diagnostic health services; optimised human mobility and smart cities, e.g. integrating mobility mining and constraint programming; and, natural resource management, sometimes referred to as computational sustainability.

3 Technical Challenges

In this talk a set of specific technical challenges are presented, motivated by the complexities of decision making in data-rich domains such as those highlighted above. There is the general challenge of integrating CP with other technical disciplines to provide a holistic solution to specific classes of problems, or to address the requirements of particular application domains. Therefore, CP must integrate with a variety of other technical domains in order to meet these challenges such as: machine learning; data mining; game theory; simulation; knowledge compilation; visualization; control theory; engineering; medicine and health; bioscience; and mathematics. Domain-specific integrations must also emerge in areas such as: life sciences, sustainability, energy efficiency, the web, social sciences, and finance. A sample of such opportunities are discussed in this talk.

References

1. Barry O'Sullivan. *Opportunities and Challenges for Constraint Programming.* Proceedings of AAAI, 2012.
2. Laurent Perron, Paul Shaw, and Vincent Furnon. Propagation guided large neighborhood search. In Mark Wallace, editor, *CP*, volume 3258 of *Lecture Notes in Computer Science*, pages 468–481. Springer, 2004.
3. Francesca Rossi, Peter van Beek, and Toby Walsh. *Handbook of Constraint Programming.* Foundations of Artificial Intelligence. Elsevier, New York, NY, USA, 2006.
4. Pascal Van Hentenryck and Russell Bent. *Online Stochastic Combinatorial Optimization.* MIT Press, 2006.
5. Gérard Verfaillie and Narendra Jussien. Constraint solving in uncertain and dynamic environments: A survey. *Constraints*, 10(3):253–281, 2005.