A CSP Approach to IT Service Management

G. Dreo Rodosek, T. Kaiser
Leibniz Supercomputing Center
Barer Str. 21
80333 Munich
Germany
{*dreo, kaiser*}*@lrz.de*

R. Rodosek
Center for Planning and Resource Allocation
Imperial College
London SW7-2AZ
England
r.rodosek@icparc.ic.ac.uk

An important task of IT service management with respect to the service life cycle is to plan the provision of services in a management environment. The goal is to provide *services* to *customers* with a high quality of service by optimizing the use of *resources* necessary for the realization of the services. IT service managers are constantly confronted with "look-ahead" problems, like how many new customers can be supported with existing resources. Such problems can be effectively and efficiently solved with a constraint-based approach, which has four key characteristics. (i) It can describe the management environment (i.e., build a model) in an incremental way. (ii) It can deal with incomplete data. (iii) It can cope with changes in highly dynamic environments. (iv) It can find optimal solutions based upon the established model.

Problems in planning IT services can be regarded as Constraint Satisfaction Problems (CSPs). A CSP is defined by a set of variables, a set of values for variables, and a set of constraints over variables. A solution to the CSP is an assignment of variables (i.e., a value for each variable) such that all constraints are satisfied.

From the modelling viewpoint, the CSP model defines all solutions to the problem whereas incomplete problem data represent a missed constraint for which a larger set of solutions is obtained, and changes in the management environment define modified constraints. The quality of solutions depends on how good the CSP model represents the problem.

From the viewpoint of solving, backtracking is a simple algorithm that solves CSP problems. However, this algorithm suffers from a lack of efficiency. The aim is to use other algorithms to solve the problem faster. These can be achieved by (i) developing a dedicated algorithm for a given problem or (ii) decomposing the problem into subproblems and solving them with well-known algorithms.

We sketch the proposed approach for the IP service. One of the important questions herewith is to recognize how many new customers (i.e., customer service requests) can be handled with the existing network infrastructure of a provider. There are two dimensions of this task, one focussing on the recognition of the existing service requests as data flows through the network and the other relating to the spare bandwidth capacity as it might be used in the future. Both dimensions can be handled with the CSP approach.

The first problem of recognizing data flows can be modelled with constraints on the following variables: (i) topology (including nodes, links, capacities of links), (ii) data traffic (acquired from network management platforms), and (iii) traffic flows. The first two variables are constants in the model while the traffic flows represent the problem variables. The input data of the model are essential for the quality of recognized data flows. By using some data traffic (e.g., ifIn/OutOctets), the model restricts ranges of each traffic flow. These ranges can be narrowed by enhancing the model with constraints over other data traffic such as information about packets and protocols. To derive the ranges of data flows, the problem is decomposed into subproblems where the paths among nodes are not changed by the OSPF routing algorithm. By doing this, every subproblem contains only linear constraints over problem variables and the simplex algorithm can solve the problem. Thus, simplex is performed twice for each variable, once to derive the minimum and once to derive the maximum value of the range. Afterwards, the solutions of the subproblems are composed by summing the ranges of derived flows, and hence recognizing data flows for the existing service requests. An example of a recognized flow is that 60%-75% of all traffic data generated by a customer is transported to the USA gateway.

Empirical results have shown very good results that can be obtained for meshed-liked topologies and unbalanced data traffic. The constraints for this class of problems define solutions with very narrowed ranges of problem variables. To derive the same quality of results for tree-like topologies, the CSP model has to contain more information about data traffic.

The second problem of identifying the maximum number of new customers can already be solved with the same model as developed for the first problem. The constraints among variables are the same for both problems, the difference is only in variables. Instead of searching for the size of data flows on each path through the network, as we did for the first problem, the variables of the second one are the maximum number of service requests.

To summarize, the strengths of the CSP approach for planning IT services are: (i) the same model can be used for different queries (i.e., look-ahead problems), (ii) it is possible to enhance the model by adding new constraints in an incremental way, (iii) efficient solution algorithms can be defined for each problem. Classical approaches such as mathematical models and simulation tools, based on discrete-event simulations or queuing theory, require much more input data and development effort to obtain a solution. These approaches suffer from a lack of robustness when faced with incomplete data, an inability to build a model in an incremental way and limited potential to cope with changes in highly dynamic environments.