A problem-adjusted genetic algorithm for flexibility design

Michael Schneider [a,*], Jörn Grahl [b], David Francas [c], Daniele Vigo [d]

[a] Chair of Business Information Systems and Operations Research, University of Kaiserslautern, Germany
[b] Department of Information Systems and Business Administration, University of Mainz, Germany
[c] Camelot Management Consultants AG, Germany
[d] Department of Electronics, Computer Science and Systems, University of Bologna, Italy

ABSTRACT

Many present markets for goods and services have highly volatile demand due to short life cycles and strong competition in saturated environments. Determination of capacity levels is difficult because capacities often need to be set long before demand realizes. In order to avoid capacity-demand mismatches, operations managers employ mix-flexible resources which allow them to shift excess demands to unused capacities. The Flexibility Design Problem (FDP) models the decision on the optimal configuration of a flexible (manufacturing) network. FDP is a difficult stochastic optimization problem, for which traditional exact approaches are not able to solve but the smallest instances in reasonable time. We develop a Flexibility Design Genetic Algorithm (FGA) that exploits qualitative insights into the structure of good solutions, such as the well-established chaining principle, to enhance its performance. FGA is compared to a commercial solver, a simple GA, and a Simulated Annealing local search on instances of up to 15 demand types and resources. Experimental evidence shows that the proposed approach outperforms the competing methods with respect to both computing time and solution quality.

1. Introduction

Demand-capacity mismatches occur in a wide range of operations like workforce planning, strategic network planning or the design of queueing systems. They are mainly triggered by uncertainty about the demand for workforce, goods and services. Demand uncertainty is driven by shortening product and service life cycles, high competition on saturated markets and the increase of product and service variety. Resource capacity levels are usually fixed for a significant time period and must therefore be set long before demand realizes. Capacity levels should match demand closely to avoid capacity shortage (when demand exceeds capacity) or idle time (when capacity exceeds demand). The use of flexible resources is an important means to counter such mismatches. Bertrand (2003) discusses several types of flexibility that have been introduced in science and practice. We consider mix flexibility, which is the ability of a resource to process several demand types without incurring high transition penalties (Koste and Malhorta, 1999).

The core idea of using mix-flexible resources is to assign several demand types to a single resource, so that the capacity

of the resource can be used for more than one demand type. Flexibility decisions can be formalized by means of a bipartite graph, where one set of nodes represents the demand types and the other set the resources. An arc between a demand type and a resource is called a *link* and indicates the capability of the resource to process the demand type. If a resource is only connected to a single demand type, it is inflexible, whereas if it has more than one linked demand type, it is mix-flexible.

Fig. 1 shows a simple example with two resources *A* and *B* and two demand types 1 and 2. Both resources have a capacity of 100 units and demand is 150 units for demand type 1 and 50 units for demand type 2. If both resources are inflexible and resource *A* can only process demand type 1 and *B* demand type 2 (as indicated by the solid links), we have a capacity shortage of 50 units at resource *A* and 50 units of unused capacity at resource *B*. Now, if resource *B* is mix-flexible and can process both demand types (by adding the dashed link), demand types 1 and 2 are linked by resource *B*. In this way, the 50 units of excess demand at resource *A* can be shifted to resource *B*. We end up with no capacity shortage and two completely utilized resources. Note also that, as illustrated in Fig. 1, in the mix-flexible case one demand type may link more than one resource in the sense that the resources are both able to process the demand type.

Originally, this generic principle has been proposed by Jordan and Graves (1995) in the context of designing flexible production networks. Since then, it has been successfully applied in several

* Corresponding author.
  E-mail addresses: schneider@bisor.de, drecore@gmail.com (M. Schneider), grahl@uni-mainz.de (J. Grahl), francas@bwl.uni-mannheim.de (D. Francas), daniele.vigo@unibo.it (D. Vigo).
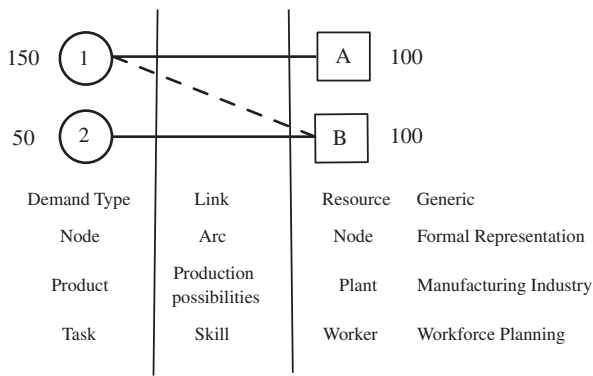
Fig. 1. Example of the benefits of mix flexibility with two resources {A,B} having capacity 100 and two demand types {1,2} with demands 150 and 50. Without flexibility, we end up with a capacity shortage of 50 units at resource A and 50 units of unused capacity at resource B. With flexibility (dashed line), there is no capacity shortage and two completely utilized resources. The figure also lists domain specific terms for demand types and resources.

practical applications in the same domain (see, e.g., Hallgren and Olhager, 2009; Kauder and Meyr, 2009) or in other industrial areas like the cross-training of workers at production lines and call centers (see, e.g., Brusco and Johns, 1998; Hopp et al., 2004; Wallace and Whitt, 2005; Nembhard, 2007) and the design of queueing systems (see, e.g., Gurumurthi and Benjaafar, 2004; Andradóttir et al., 2007). Fig. 1 lists the appropriate domain-specific terms for the two major application areas. However, as the contributions of our paper are not restricted to an application domain, we stick to the generic terms for the remainder of the text.

The challenge common to the different application areas is to find the most profitable assignment of demand types to resources when confronted with stochastic demands. This decision is modeled by the Flexibility Design Problem (FDP). The FDP determines the optimal assignment of demand types to resources such that the total expected profit is maximized. We formulate the FDP as a two-stage stochastic program, however, the FDP is NP-hard even for deterministic demands (Garey and Johnson, 1979). In our numerical studies, we observe the inability of standard solvers to solve all but the smallest FDP instances to optimality in a reasonable amount of time. Motivated by this, we develop a Genetic Algorithm (GA) specifically tailored to FDP, denoted as Flexibility Design GA (FGA). FGA incorporates problem-specific knowledge on the structure of high-quality FDP solutions to guide the search. Specifically, it uses the flexibility principles of Jordan and Graves (1995), which have proven to be a reasonable and robust strategy for designing flexible networks. We thus demonstrate how qualitative principles can be incorporated into a quantitative solution approach to increase computational efficiency.

We present a computational study on FDPs with up to 15 demand types, 15 resources and 500 demand scenarios. To the best of our knowledge, no benchmark instances for the FDP or a closely related problem are available and no efficient metaheuristic solution method for the FDP has been proposed. Therefore, we generate several sets of random test instances. As comparison methods for our approach, we use the commercial solver Xpress-Optimizer (see Fico Xpress Optimization Suite, 2011), a Simple GA (SGA) and a Simulated Annealing (SA) approach. FGA demonstrates superior performance, consistently finding flexibility configurations with higher expected profits in a fraction of the time required by the comparison methods.

The paper is organized as follows. Section 2 reviews the literature and describes the flexibility principles developed by Jordan and Graves (1995). A formulation of the FDP as a two-stage

stochastic program is given in Section 3. Section 4 describes in detail the FGA for solving this problem. Experimental results are presented in Section 5. The experiments evaluate the run-time and solution quality of the FGA and the three comparison methods. Moreover, we study the effectiveness of each problem-specific extension that we integrate into FGA. Section 6 concludes the paper and gives a brief outlook on future research.

2. Literature review

Jordan and Graves (1995) present a pioneering work on the structure of high-quality solutions for the design of flexible networks. They study the optimal assignment of products to plants in a simplified multi-product and multi-plant manufacturing network. In simulation studies, the flexibility benefits of different configurations are compared using expected sales and expected capacity utilization as performance measures. Central to their work is the concept of a chain. A chain is defined as a group of products and plants which are all connected, directly or indirectly, by product assignment decisions. In other words, the bipartite graph representation of a chain is connected, i.e., it consists of one connected component.

Jordan and Graves (1995) propose the following flexibility principles, which have heavily influenced academic research and industry practice since then: (1) Chains which encompass as many plants and products as possible are superior to configurations that contain several chains. (2) A complete chain which has all nodes of the graph in a cycle has virtually the same benefits as a totally flexible network, i.e., a network in which each product is connected to all resources. From a flexibility viewpoint, adding further links beyond a complete chain configuration is almost useless. (3) In general, there are several possibilities to configure a production network in a way that optimizes flexibility benefits. (4) When adding flexibility to an inflexible network, the demand for each product should be balanced with the capacity of the plants the product is directly connected to. Analogously, the capacity of each plant should be balanced with the demand of the products the plant is directly connected to.

Iravani et al. (2005) develop domain-independent flexibility measures based on the number of possibilities to shift demand in a network. They analyze different structures for mix flexibility and develop indices to quantify a system's ability to respond to variability in its environment. Their results underline the power of chain-like flexibility configurations. Aksin and Karaesmen (2007) address mix flexibility in systems with parallel flow using a graph–theoretic approach to determine the maximum achievable throughput under a particular network structure. Their results on structural properties of mix-flexible configurations also support the concept of chaining. Bassamboo et al. (2010) assess the problem of capacity and technology selection based on the newsvendor network model of Van Mieghem and Rudi (2002) to deduct optimal flexibility configurations. In their evaluation, the optimal solutions are not chain structures, yet chain-based solutions have only small optimality gaps. Chou et al. (2010) provide analytical justification for the efficiency of chaining structures. They analyze the performance of chain structures and identify a class of conditions under which chain structures yield performances similar to those of fully-flexible systems.

A strand of the related literature deals with the model-based strategic planning of a company's production network. Apart from mix flexibility decisions, these approaches typically consider additional issues like capacity choice, detailed operational work shift models (Bihlmaier et al., 2009), raw material supplier selection (Alonso-Ayuso et al., 2003), net present value based optimization (Fleischmann et al., 2006) and uncertain demand,