High-Level Modelling and Solving for Online, Real-Time, and Multiagent Combinatorial Optimisation

Alexander Ek

Combinatorial problems are notoriously hard decision-making problems to solve and appear in many fields and industries. Examples include vehicle routing in logistics, scheduling in manufacturing, resource allocation in finance and humanitarian efforts, rostering staff in health institutions, and control systems in automation endeavours. The field of combinatorial optimisation develops techniques that help find quality solutions to these kinds of problems. The importance of finding high-quality solutions to these problems has resulted in impressive advances in this field, including solving technologies such as CP, MIP, and SAT, as well as modelling languages such as Mini-Zinc, Essence, OPL, and AMPL.

Many modern combinatorial problems are *dynamic*, i.e., new information continually arrives, and the solution must repeatedly be refined to update existing decisions and make new ones if needed. Typically, building a system to solve these problems requires substantial ad hoc coding, despite commonalities therebetween. Unfortunately, combining dynamism and combinatorial optimisation reveals a gap in the literature that precludes using the modelling and solving technologies mentioned above.

This thesis helps fill this gap by presenting a novel framework for modelling and solving dynamic combinatorial problems that avoids ad hoc coding by taking advantage of and extending previous advances in modelling languages. The framework allows modellers to declaratively specify the dynamism of the problem being solved, which is then used to solve the problem with online and real-time optimisation automatically. This framework is then enhanced with a systematic *garbage collection* mechanism, which ensures that the internal representation of the evolving and growing problem remains small, even as time progresses. This is complex because such a mechanism must account for how past decisions affect future ones. Without garbage collection, internal representations quickly get too large, impeding compilation and solving time. The enhanced framework is implemented and evaluated empirically.

Many combinatorial problems concern multiple stakeholders with conflicting demands, making fairness and balance essential. Substantial work exists in adjacent

areas, such as economics, social choice theory, game theory, and multi-objective optimisation. However, to the best of our knowledge, the intersection of combinatorial optimisation, fairness, and dynamism has never been studied. Because of the multifaceted complexity, formal guarantees on fairness are hard to give, and empirical experimentation becomes necessary. This thesis presents a new framework (in which the aforementioned framework can be used) capable of addressing the issue of fairness in dynamic combinatorial problems, and which allows easy experimentation with different ideas of fairness. The framework is founded upon an adaptive priority rebalancing principle that gives unjustly treated stakeholders higher priority in future decision-making. The results show that the framework can easily be used to achieve remarkable increases in fairness without significant loss in overall resource utilisation efficiency.

Even without dynamism, there is a significant gap in combinatorial optimisation with fairness considerations of inequality. This thesis presents two novel constraint propagation algorithms for two inequality metrics: variance and the Gini coefficient. They can produce explanations required for clause learning, making them usable in many state-of-the-art solving methods. These algorithms allow speedups in solving problems in which using these two inequality metrics is desirable.

School: Monash University, Clayton VIC 3800, Australia Supervisors: Guido Tack Maria Garcia de la Banda Andreas Schutt Peter J. Stuckey Graduated: Friday, 23 September, 2022 Link to full text: https://doi.org/1.26180/21588144.v1

2