Dynamic Contracting in Infrastructures (PhD Dissertation Abstract)

Joris Scharpff
Delft University of Technology
Deparment of Computer Science, Algorithmics Group
Mekelweg 4, HB 07.080
2628 CD Delft, The Netherlands
j.c.d.scharpff@tudelft.nl

Abstract

Contracting for infrastructural maintenance has changed significantly over the last couple of decades from the well-known 'regulatory' contract towards more innovative and flexible agreements. Indeed these newer methods offer promising benefits: increased flexibility, better performance and consequently less cost. Nonetheless, these benefits are accompanied by higher levels of uncertainty and more possibilities for opportunistic behaviour on a network level. In order to tackle these additional complexities, more dynamic and flexible contracting methods are required.

In the combined research 'Dynamic Contracting in Infrastructures' we propose a two-phase approach for the procurement and realisation of maintenance contracts in infrastructures, consisting of a procurement phase and an execution phase. In the procurement phase maintenance activities are identified, priced and auctioned to third-party contractors. In the execution phase we focus on the planning and execution of maintenance activities using a novel dynamic incentive mechanism that stimulates contractors to implicitly consider contracted demands and collaborate on a network level.

The main contributions of this PhD is an innovative dynamic payment mechanism for the executional phase of the contracting procedure we propose. Using a novel combination of mechanism design and planning under uncertainty we implicitly stimulate contractors to consider contracted objectives, other than monetary, and collaborate on a network scale, even when the execution of maintenance is uncertain.

1 Introduction

In the past two or three decades we have witnessed a considerable change in the procurement and realisation of infrastructural maintenance projects. The introduction of Public-Private Partnerships (PPP) in the late 70s and early 80s has resulted in more innovative, risk-sharing contract forms, adopted by governments and public institutions world-wide (Altamirano 2010; Behn and Kant 1999; Chi, Arnold, and Perkins 2003; Tieva and Junnonen 2009). Indeed such contracts are expected to offer various benefits over the more 'classical' regulatory contract[1]: increased flexibility, more innovation, better performance and subsequently lower costs (Altamirano 2010).

These benefits, however, are accompanied by a higher level of uncertainty and introduce additional possibilities for opportunistic behaviour. Particularly in the long-term, performance based contracts – employed in infrastructural maintenance – these undesired effects are likely to arise, resulting in unsatisfactory results or even total failures (O'Hare, Leone, and Zegans 1990).

Moreover, performance based contracts allow for a greater degree of freedom in project implementation although consequences thereof are commonly neglected. In most scenarios, there is a misalignment in the objectives of both parties: public institutions seek to optimise social welfare whereas contractors are focussed only on profit.[2] These different interests frequently give rise to conflict, a problem identified already in the work by Pigou in the beginning of the 20th century (Pigou 1920). Still, this *social cost* (Coase 1960) caused by the contractors in optimising their personal objectives is often not adequately accounted for in the contract.

One reason for the incorporation of social cost into the contract is to make the contractor aware of the consequences his choices have on society. Actions that hurt society more should cost the contractor more money. This allows the contractor to make its own trade-offs in deciding what action to take. Secondly, we can introduce sharing of social cost to stimulate contractor co-operation. This is definitely an opportunity that is not present in current contracting procedures.

Consider for instance two contractors performing maintenance within the same region, both servicing a different a different segment of the road network. If we express social cost in terms of the additional congestion caused by maintenance then the profit of a contractor is influenced by the planning choices of the other. When both contractors decide to plan their work on the same date, congestion will increase enormously and therefore both contractors are charged a higher social cost payment than they would have been if they chose two disjoint periods to do the work.

[1] We have adopted the term regulatory contract from (Behn and Kant 1999) to denote all the contract approaches that are based on controlling contractors by regulation.

[2] One could argue that contractors are focussed on continuation instead, however in our game theoretical setting we assume utilitarian agents.

This dependency can be exploited to stimulate more (social) cost-effective decisions and collaboration. In the example given above, the contractors could negotiate and make sure that their work plans are disjoint in order to reduce their social cost and hence increase their profit. On the other hand, if the accumulated social cost of performing the work separately is larger than when both contractors choose the same period, that outcome should be stimulated by sharing social costs. Intuitively we seek to find the a payment structure for social cost that stimulates contractors to collaborate on finding the combined plan that has the lowest accumulated social cost.

Such a new approach to maintenance contracting requires a redesign of current procedures. In this research we focus on the development of a novel dynamic contracting framework that tackles these problems. Moreover, we study the applicability of such a procedure on real-world problems, from both a theoretical perspective, combining state-of-the art literature into a new mechanism, as well as a practical perspective, through the use of expert validation and serious gaming.

The remainder of this dissertation is organised as follows: in Section 2 we present the dynamic contracting procedure. As this PhD work will be focussed on the executional phase, Section 3 and Section 4 respectively discuss the research questions and the solutions we propose for this phase. Section 5 discusses future work.

2 Dynamic Contracting

The previous section has made clear that although more innovative contract forms are very promising, there are still significant problems to deal with. The main issues in current contracts can be related to their more static nature; contracts are commonly fixed at the beginning of a project and are not sufficiently capable of adapting to changes. Especially in long-term projects, where a lot of (unforeseen) changes occur, current contracts are hardly adequate to ensure a successful result. Moreover, the network aspect of infrastructural maintenance is commonly neglected in current contracting procedures, thereby failing to work towards system optimal asset management.

In overcoming these issues, we propose a two-stage, network-based dynamic contracting procedure (depicted in Figure 1, see also (Volker et al. 2012)) that:

- integrates the network and the network users as central components in the tendering and execution of the contract,
- provides flexibility and robustness for long-term contracting,
- incentivises contractors to achieve high levels of infrastructural quality and
- increases contractor autonomy while stimulating collaboration and requiring less in-house knowledge at the asset owner or manager.

In the procurement phase we design the framework within which maintenance is to be performed. Based on the project

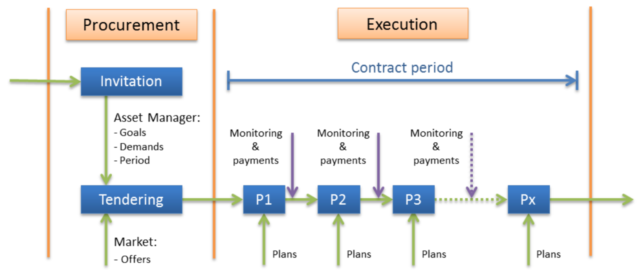

Figure 1: Our proposed dynamic contracting procedure. In the procurement phase maintenance activities are identified, priced and subsequently tendered to contractors. In the second phase these activities are planned and executed in a dynamic fashion.

goals and demands, specified by the asset manager, the network is divided in segments which are put up for tender. In addition a pricing scheme is announced that captures the social cost of maintenance, as discussed in the introduction. The market i.e. contractors can submit offers for these segments of the road network, basing their prices on the expected cost of maintenance (both their private and social costs) given the asset demands. As a result of the first phase we have a group of contractors, each responsible for a part of the infrastructure, a set of asset demands and a pricing mechanism that corresponds to the offered price, incorporating the social cost charges. Note that in this phase no actual plans for maintenance are developed.

The resulting frame contract from the procurement phase is used to define the boundaries of the execution phase in which the actual planning and execution is performed. It is now up to the contractors to identify the maintenance activities that should be performed on their own part of the network and develop a joint, socially (near) optimal schedule for these activities.

As we are dealing with long-term contracting in a contingent environment, we use a periodical planning approach. Using one plan for the entire contract duration is unrealistic; infrastructural maintenance is vulnerable to unexpected delays, possibly affecting the entire schedule. Instead we periodically develop plans for a shorter period of time, e.g. one month.

The Asset User

Although many researchers and practitioners underline the importance of incorporating the asset user in the asset management process, this is rarely the case in asset management literature. In the model we propose, we represent the asset user by including the social cost of maintenance in our method. The social cost represents the monetary value of the total (negative) impact of a joint maintenance plan. Hence, the social cost depends on all the scheduled maintenance activities over the entire network.

Service provides are charged payments relative to their share of the social cost. For instance, in the case of road maintenance, causing more congestion means a larger social cost payment. In this way, contractors are implicitly stimulated to plan their work at favourable times from the asset users point of view, as blocking roads during rush hours is more costly than for example doing the same at night.

An additional major advantage of the social cost payments is that it also implicitly creates dependence between contractors, as social cost is computed over the joint maintenance plan. Recall the two-contractor example from the introduction, where two contractors can either decide to perform their maintenance on the same day or chose different times. Using our social cost as an incentive, we align the most profitable outcome for the individual agents with the socially optimal outcome.

3 Execution Phase

Under performance based contracting, contractors are given a lot more freedom as opposed to regulatory contracting: only goals are specified, leaving the implementation of these goals up to the contractor. However, the main focus of a contractor is to maximise his profit and therefore we require (monetary) incentives to stimulate optimisation of other goals such as asset quality or network throughput. We want contractors to be autonomous but in their choices they should also regard the (financial) impact of these choices. Based on our preliminary survey we have identified the research questions presented below.

Q1: How can we model the problem of maintenance planning for competitive agents such that one the one hand we optimise global objectives and, on the other, optimise individual contractor profits?

In infrastructural maintenance we often have a single public road authority responsible for the quality, throughput and costs of the network, but the actual maintenance is performed by commercial, third-party contractors, typically interested primarily in maximising their profits. Road authorities face the problem of aligning objectives; we introduce (monetary) incentives for the service providers to consider the global objectives. But an agent servicing one part of the network also influences agents in other parts as his work has a negative impact on the traffic flow. As a consequence, the payments we introduce lead to very high throughput penalties for all agents if they do not coordinate their maintenance plans.

In this PhD we aim to incentivise contractors to implicitly account for these (contracted) global objectives. Our main contribution is the application of a combination of stochastic planning and dynamic mechanism design to realise coordination between non-cooperative agents. Typical one-shot mechanisms often studied in mechanism design are not suitable for contingent and repeated settings. Instead we focus on dynamic mechanisms that define payments over all expected outcomes such that in expectation it is in the agents best interest to be truthful during the entire plan period.

Sub-questions

- Maintenance planning is inherently a multi-objective problem (e.g. costs, quality state, throuput, durability, . . .). How can we adapt our mechanism to deal with multi-objective scores and payments such that we can still incentivise agents to consider these objectives and report truthful?

Q2: Given a mechanism such that contractors indeed account for global objectives, how can they plan their activities on a network level in such a way that they are cost-optimal in the autonomous, multi-agent setting?

Even if a contractors wants to keep the social cost to a minimum, it is only one player in a multi-agent scheduling problem and hence it is affected by the choices other contractors make. Considering for example the total additional traffic generated by maintenance as the basis for our social cost function, maintenance activities on the network might influence other agents as well. So by solving its scheduling problem individually, contractors will most likely not obtain the most profitable maintenance plans.

We study first a centralised approach in which the asset manager (road authority) is given the task to find *efficient* joint plans. Efficient here refers to the plans that are socially optimal, i.e. no improvement for one agent can be made without harming another agent. Another approach, one that we also implement as the main method in a *serious game* (see Section 4), is a myopic approach analogue to the best-response planning proposed by Jonsson in (Jonsson and Rovatsos 2011).

For the centralised planning, we require a mechanism such that the asset manager is able to elicit contractor maintenance costs. However, contractors might not be willing to share such vital business information concerning their costs. Even if contractors are willing to share, they might report false information to the asset manager in order to increase their profit. Such *opportunistic behaviour* should be countered in a centralised approach, something we want to achieve through mechanism design. Note that the application of such a mechanism requires optimal solutions to base payments on, otherwise truthfulness is not necessarily guaranteed. As a consequence, we cannot directly apply standard approximation techniques to the problem as it might invalidate the mechanism.

Sub-questions

- Can we exploit the structure of the planning problem to decrease the run time required for finding (near-)optimal solutions centrally?
- How can we relax the optimality requirement of mechanisms while still preserving the mechanism desiderata? Moreover, which conditions are sufficient (or preferably necessary) for 1) truthfulness, 2) individual rationality and 3) budget balance in dynamic mechanisms using approximation? Can these properties be generalised into a class of dynamic approximation mechanisms?

Q3: Given the mechanism that results from questions 1 and 2, what is are the trade-offs of using a decentralised approach in comparison to centralised method?

Developing joint maintenance plans centrally has the main advantage that we can find efficient (i.e. socially optimal) joint plans, hence the best solution for the contractors, the asset manager and the asset users combined. However, centralised solving has two major drawbacks. First, contractors have private information about their maintenance costs, which have to be elicited by the asset manager in order to develop optimal plans. Second, assuming we are able to obtain this information from the contractors, finding optimal joint plans poses a complex optimisation problem.

In a decentralised method, contractors need only to submit plans instead of information about their costs. Using a Nash-like iterative improvement of the joint plan, contractors gradually improve on a joint plan. This method does not suffer from privacy issues and is computationally much more feasible, however at the cost of (expected) lower joint plan quality. Instead of developing an optimal joint plan, contractors work towards equilibria that are not necessarily guaranteed to be (near-)optimal. Still, this method can be easily translated to practical settings.

Sub-questions:

- Given that all players use a best-response strategy to improve on the joint plan, does an equilibrium always exist? If not, can we find characteristics of the problem that guarantee the existence of (and convergence to) an equilibrium? Can we derive any guarantees on the quality of such equilibria?

Q4: Given that execution of maintenance activities is inherently contingent, how can we design our mechanism such that its desiderata still hold under uncertainty?

In the ideal setting we can use a simple one-shot mechanism to find jointly optimal plans in advance and subsequently execute them without having to interfere. However, performing maintenance is inherently contingent and uncertainties arise in many different aspects. The asset state might be unknown, maintenance operations might take more time than planned for due to e.g. bad weather conditions, and so on.

This uncertainty must be taken into account while developing plans and reflected in the payment mechanism.

Sub-questions

- How can we extend current results on multi-objective mechanism design to the dynamic setting?

4 Plan Coordination under Uncertainty

The research questions presented in the previous section have directed our research towards a novel combination of mechanism design with stochastic planning. Instead of a single plan, contractors develop policies that dictate the best action to perform *given the current state* the contractor is in. The mechanism payments are modified accordingly: rather than using one fixed payment for a plan, we use dynamic payments such that contractors *in expectation* make the most profit when they consider global goals.

In (Scharpff et al. 2013), accepted for presentation at the ICAPS 2013 conference, we discuss a real-world application of our dynamic mechanism and planning method for coordination under uncertainty. Basically, contractors are rewarded for their expected contribution towards global goals such as road quality and maximal network throughput, but have to balance this with their own maintenance revenues and costs.

We use a dynamic variant of the Vickrey-Clarke-Groves (VCG) ((Vickrey 1961; Clarke 1971; Groves 1973)) type mechanism to determine the payments each contractor pays or receives for his participation. Informally the VCG-mechanism charges each contractor the harm they cause to other agents (including the asset manager and asset user) by their presence, i.e. their maintenance on the network. Using this mechanism, we strive to develop optimal joint maintenance plans in this work. Our evaluation however shows that this is computationally hard and therefore more research is required in this area.

Serious Game

Next to our theoretical work on mechanism design and planning, we also concentrate on the practical aspects and impact of our dynamic contracting procedure. First of all, the maintenance planning problem, subject of study in (Scharpff et al. 2013), has been obtained through interviews and discussions with domain experts. Moreover, in order to validate our theoretical work, we are developing a serious game that tests the concept of dynamic contracting (to be also submitted to the ICAPS Application Showcase). The major goals of this game are:

- Studying whether our novel contracting method can be used in practical scenarios, and whether practitioners are likely to accept and adopt our method.
- Creating awareness and support amongst practitioners regarding the impact of (coordinating) maintenance activities on a network level. Using this tool we want practitioners to get a feel for our novel and progressive concept, increasing the likelihood of acceptance.
- Validation of the payment mechanism. Human players will most likely not be perfectly rational, therefore we study the strategies played by human planners and the resulting outcomes.
- Closing the gap between theoretical concept and realistic contracting. This will increase the likelihood of practical implications.

In this game, players take on the roles of contractors and have to maximise their profit over a given portfolio of maintenance activities (see Figure 2). They are supported by an automated planner that provide insight into payments and costs, and is able to provide plan suggestions, based upon the work in (Scharpff et al. 2013).

5 Future Work

Our current research activities are currently divided over two tracks: 1) finding (more) tractable methods to solve the maintenance planning problem and 2) the development and improvement of our serious game through experimental sessions. For the first track, we are now considering using state-of-the-art RDDL planners as the core solver for the problem, but also we try alleviate the complexity through exploitation of the problem structure.

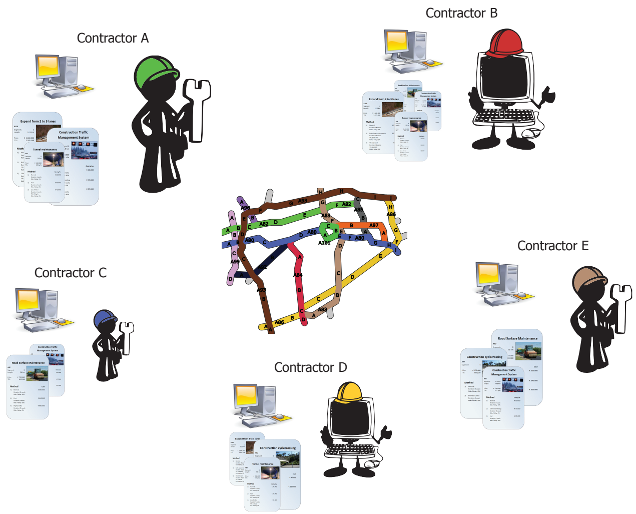

Figure 2: In the game, contractors (played by humans or computer agents) need to plan their given portfolio of activities on the network in the most profitable way. Their portfolios are represented by task cards, that specify the details of each activity.

In future work we want to study approximation methods for the maintenance planning problem. The main difficulty is that while approximation makes finding maintenance plans easier, the payment mechanism does not have to be incentive compatible anymore (i.e. truthful). The payment mechanism has to be adapted to remain its desiderata, but this depends largely on the used approximation technique.

Another extension we are addressing in future work is that of multi-objective planning. Maintenance planning is inherently multi-objective: for example, one cannot simply compare one euro of maintenance cost to a quality improvement of 2%. Currently we operationalise all objectives other than cost into monetary values using rewards and fines. Using pure multi-objective approaches would be more realistic, but makes both the planning problem and mechanism design problem much more complex.

References

Altamirano, M. 2010. *Innovative contracting practices in the road sector: Cross-national lessons in dealing with opportunistic behaviour*. Ph.D. Dissertation, Delft University of Technology.

Behn, R., and Kant, P. 1999. Strategies for avoiding the pitfalls of performance contracting. *Public Productivity & Management Review* 22(4):470–489.

Chi, K.; Arnold, K.; and Perkins, H. 2003. Privatization in state government: trends and issues. *SPECTRUM-LEXINGTON* 76(4):12–21.

Clarke, E. 1971. Multipart pricing of public goods. *Public choice* 11(1):17–33.

Coase, R. 1960. Problem of social cost, the. *Journal of Law & Economics* 3:1.

Groves, T. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* 617–631.

Jonsson, A., and Rovatsos, M. 2011. Scaling up multiagent planning: A best-response approach. In *Proceedings of the Twenty-First International Conference on Automated Planning and Scheduling*, 114–121.

O'Hare, M.; Leone, R.; and Zegans, M. 1990. The privatization of imprisonment: A managerial perspective. *Private prisons and the public interest*.

Pigou, A. 1920. *The Economics Of Welfare*. Macmillan and Co., London.

Scharpff, J.; Spaan, M.; Volker, L.; and De Weerdt, M. 2013. Planning under uncertainty for coordinating infrastructural maintenance. In *Proc. of the international Conference on Automated Planning and Scheduling*. To appear.

Tieva, A., and Junnonen, J. 2009. Proactive contracting in finnish ppp projects. *International Journal of Strategic Property Management* 13(3):219–228.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16(1):8–37.

Volker, L.; Scharpff, J.; De Weerdt, M.; and Herder, P. 2012. Designing a dynamic network based approach for asset management activities. In *Proc. of 28th Annual Conference of Association of Researchers in Construction Management (ARCOM)*.