Road Pricing

Kai Nagel

15.1   Basic Information

Entry point to documentation:

`http://matsim.org/extensions` → roadpricing

Invoking the module:

`http://matsim.org/javadoc` → roadpricing → `RunRoadPricingExample` class

Selected publications:

Rieser et al. (2007a, 2008); Grether et al. (2008)

15.2   Introduction

Roadpricing is a controversial policy measure (e.g., Button and Verhoef, 1998). Its implementation in MATSim is conceptually straightforward (Rieser et al., 2007a, 2008; Grether et al., 2008): Essentially, for each vehicle entering a link at a given time, the appropriate toll is computed and charged to the vehicle's driver. The scoring function will pick this up by the term (see Equation (3.4))

$$S_{trav, car, q} = \ldots + \beta_m \cdot \tau + \ldots ,$$

where τ is change in the monetary budget invoked by all toll payments (usually negative) and β_m is the marginal utility of money (also see Chapter 3 and Chapter 51). The driver then takes this into account making decisions, e.g., for route choice, departure time choice, mode choice, destination choice, etc., and then trades off toll payments with other elements of his or her scoring function.

It should be clear that this automatically picks up all kinds of heterogeneities, for example:

- Traveling at a different time may lead to a different toll, but possibly also to different schedule delay costs (Section 3.2.5).
- Different vehicle types may be charged different tolls (Kickhöfer and Nagel, 2013).
- Different travelers may have different time values (Nagel et al., 2014), which may even vary according to the time of day.

However, one challenge is that the innovative modules (Section 4.5) must be consistent with the scoring now modified by road pricing. The approach just described will not work if, for example, the router consistently generates toll-avoiding routes for a synthetic person with a high time value, who would normally wish to pay for a faster option. In a case like this, if a suitable route is never generated, the scoring cannot identify it, giving the choice process no chance to select it in subsequent iterations.

However, processing every detail for each individual, i.e., not only the marginal utility of money, but also specific time pressure at the route search time, is quite complex.

An alternative approach is to make the router *randomizing*, i.e., to run it with a randomly generated time value every time necessary for a given person. Computational experiments with this approach produce solutions for synthetic travelers approximately as good, or even better, than an "engineered" router (Nagel et al., 2014). At the same time, the software consistency burden is significantly reduced, noticeable in the smaller amount of information to be extracted from the agent during each router call.

15.3   Some Results

15.3.1   *Effect of an Afternoon Toll on Morning Traffic*

In a first demonstration of capabilities, an afternoon toll for the Zürich area was simulated. While this is an unlikely policy scheme, it still clearly demonstrated the advantage of the integrated approach over other approaches. Not only did the synthetic travelers switch to public transit, but they also did so for the morning rush hour, where no toll was charged (Figure 15.1). Thus, the MATSim approach proved its ability to affect the whole daily plan, not just the trip. For more information, see Rieser et al. (2008).

15.3.2   *Income-Dependent Values of Time*

Similar to Rieser et al. (2008), Kickhöfer et al. (2010); Kickhöfer (2014) introduced a distance-based morning peak toll on the same links between 6:30 am and 9 am. Toll levels were incrementally increased from 0.28 *CHF/km* up to an almost prohibitive price of 44.80 *CHF/km*. The studies assume income-dependent utility functions with a decreasing marginal utility of money. The goal was to (i) identify the welfare-maximizing (see e.g., Tirachini et al., 2014, Section 2.5) toll level, which is potentially dependent on the aggregation rule of user benefits (see Chapter 51), and (ii) to investigate distributional aspects of such pricing schemes. The studies showed that changes in travel patterns resulting from the morning peak toll impacted the whole day, affecting traffic patterns in the afternoon. Furthermore, the study showed that such a parametric approach is capable of identifying the welfare-maximizing toll level. However, results also indicated that the overall welfare effect level depends strongly on the aggregation rule for user benefits, i.e., if one first monetizes
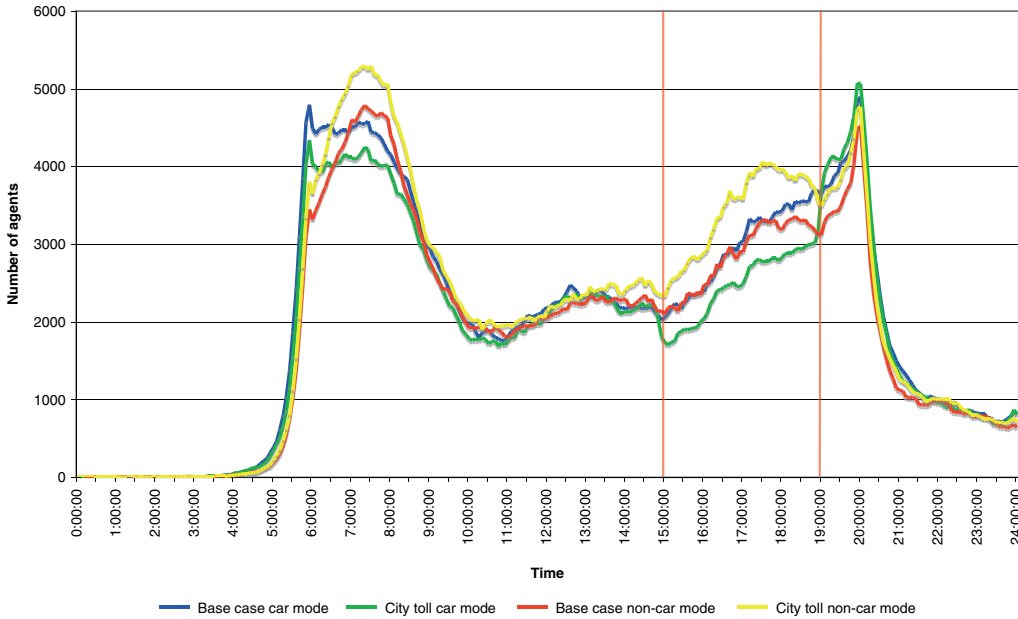
Figure 15.1: An afternoon city toll (between 3 pm and 7 pm) affects mode choice not just during the toll time, but also in the morning.

Source: Rieser et al. (2008)

individual utilities and then adds up, or first adds up utilities and then monetizes. Even the sign of that effect might not be stable depending on that choice. For more information, please refer to the two studies above.

15.3.3   Integrated Passenger and Freight Toll Simulation for the Gauteng Province in South Africa

A large scale application was undertaken for the Gauteng province in South Africa (Chapter 69). It is based on the so-called e-toll, which was switched on in December 2013. The e-toll should, logically, charge different rates for different vehicle types, with higher rates for heavy trucks. Again, logically, this should go along with higher time values of the driver-vehicle-units. Somewhat surprisingly, this turned out to be difficult to do with the MATSim software structure in place when the project was started in 2008. While it was easy to charge the freight vehicles a higher toll, it was difficult to give different replanning methods and different scoring function to the freight population; it was essentially impossible to feed the router with different time values for the freight population. This was an important driver for much development in recent years, including making the scoring function more accessible (Section 45.2.10), allowing different replanning strategies for different sub-populations (Section 4.5), and reducing consistency requirements between the router, the vehicle-based toll and the driver-based scoring function (Nagel et al., 2014).

The simulation, as expected, predicts reduced traffic volumes on the tolled roads and increased volumes elsewhere (Figure 15.2).
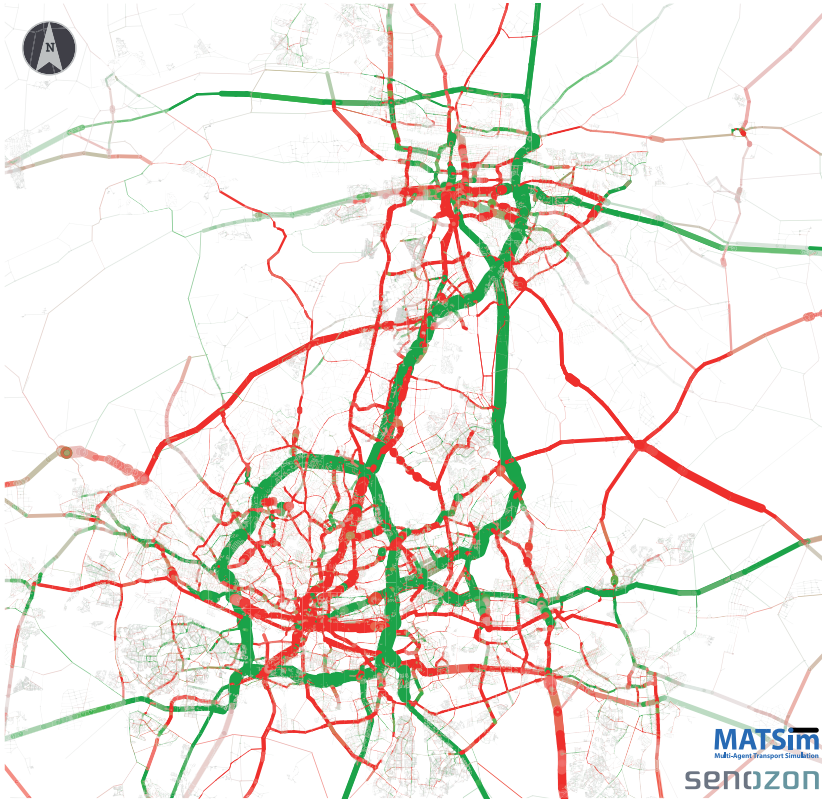
Figure 15.2: Predicted differences in link volumes after introduction of the toll (red: higher volumes, green: lower volumes).

15.4   Invocation

15.4.1   Minimal

A minimum amount of infrastructure is necessary when running roadpricing from the command line. For this, the MATSim JAR, its libraries, *and* the roadpricing JAR need to be downloaded, either from a release or from the nightly builds (Section 44.3.6). After unzipping all zip files, the necessary command is (may need slight refactoring with new formats):

```
java -Xmx2000m -cp MATSim.jar:roadpricing-.../roadpricing-
    ...jar org.matsim.roadpricing.run.RunRoadPricingExample config.xml
```

where `config.xml` needs to contain a section

```xml
<module name="roadpricing" >
...<param name="tollLinksFile" value="<path>/<tollfilename>" />
</module>
```

The toll file looks like this:

```xml
<roadpricing type="link" name="abc">
   <links>
      <link id="11">
         <cost start_time="05:00" end_time="10:00" amount="1." />
```

```
            <cost start_time="17:00" end_time="20:00" amount="1." />
        </link>
        <link id="12" />
    </links>

    <!--this is for all links with no cost entry above:-->
    <cost start_time="05:00" end_time="10:00" amount="2.00"/>

</roadpricing>
```

As one can see, there is a section where each link can be entered separately. A separate cost structure for each link is also possible. All links that are listed without a cost structure employ the general cost structure listed at the end. Links not listed are without toll.

15.4.2   Toll Schemes

Link toll   The example refers to the "link" toll scheme, indicated by type=`"link"`. It charges the amount specified on the link.

Distance toll   Another useful scheme is "distance", indicated by type=`"distance"`. Here, the amount is interpreted as amount per length unit (see Section 2.2.1). This is most useful, with only a list of tolled links and a uniform distance cost for all these links noted at the end of the file.

Area toll   The simulation of an *area toll*—i.e., a toll where one has to pay a flat fee for a given time period, often a day, once one drives anywhere inside the area—suffers from a combinatorial challenge: driving through the tolled area early in the day may only pay off if one can re-use the permit later in the day. The code, in principle, addresses that by routing the agent twice: once under the assumption of a zero toll and once under the assumption of a very large toll. Afterward, the toll is added to the generalized cost of the first option, then both options are compared. In the end, the approach suffered from the same consistency burden as the general approach (see end of Section 15.2): the router made the decision about the better variant, rather than leaving the decision to the agent. It should be re-implemented using the same principles as Nagel et al. (2014).

Cordon toll   The cordon toll scheme was derived from the area scheme; one could use the same file, listing all area links, for the cordon toll as well. The code ensured that toll was only charged when a vehicle moved from an untolled link to a tolled link—thereby effectively crossing the cordon. One difficulty with this approach: confusion ensues if there is no connected area and several links in sequence are tolled instead. Then, if these links are connected, the toll is only charged on the first of them; if there is a small section missing, perhaps overlooked, the toll is charged again.