2000

A Prototype Web-based Carpooling System

Peter Keenan
University College Dublin, peter.keenan@ucd.ie

Shane Brodie
Intel Ireland, shane.brodte@intel.com

A Prototype Web-based Carpooling System

Peter Keenan, Department of Management Information Systems, University College Dublin, Dublin 4, Ireland, *Peter.Keenan@ucd.ie*
Shane Brodie, Intel Ireland, Leixlip, Co Kildare, Ireland. *Shane.Brodie@intel.com*

Abstract

This paper discusses a web based automated system for carpooling. The design issues for such a system are discussed and a prototype system that integrates a heuristic model and a GIS is described.

Introduction

While the cities of the developed world differ in many respects, traffic congestion is one problem almost all cities have in common. This problem is particularly acute in Dublin, the largest city in Ireland. This reflects the fact that the greater Dublin region has suffered from a long-term lack of investment in transport infrastructure. Consequently, Dublin lacks either the extensive highway networks characteristic of North American urban areas or the excellent public transport facilities found in most European cities. Such public transport as does exist is largely focussed on travel to the city centre, with fewer links between suburban areas. In this situation a high proportion of commuters use their cars for their daily journeys, which has led to the chronic traffic congestion.

The traffic crisis in Dublin has been aggravated by rapid economic growth, as Ireland has performed better than other EU countries recently. Economic growth has been especially rapid since 1994 with the economy growing by more than 7% each year. This reflects a rapid increase in employment, the numbers employed increased by 6% in 1999. This also means an upsurge in the numbers commuting to work. Economic growth is reflected in new car sales, which were three times as high in 1999 as twelve years earlier. This trend is continuing, in the current year (2000) new car sales have increased by more than forty percent in the first Quarter. Historically per capita car ownership in Ireland has been relatively low, partly because cars are heavily taxed. However, rapid economic growth is bringing car ownership in Ireland towards a level associated with the most prosperous European countries. This trend still has some way to go, but it has ominous implications for future traffic congestion. The improved economic conditions have provided ample funds for the public authorities to invest in transport projects, for example the proposed light rail system, but there is inevitably a long delay between the conception and delivery of these schemes. The rapid pace of change has often served to compound this delay, as earlier plans are overtaken by events and are replaced by more ambitious proposals.

The improved economic conditions have partly resulted from the establishment in Ireland of large computer industry multinationals. These organizations are typically located in new facilities on the outskirts of Irish towns. These locations are typically difficult to reach by public transport. One of these is the Intel Corporation, which has its European headquarters at Leixlip, Co. Kildare, on the Western side of Dublin. This plant employs some five thousand people, almost all car owners. The large numbers travelling to the Intel plant are sufficient to cause serious traffic congestion in the locality, despite road improvements in recent years. Intel was faced with significant expenditure in building car parks and sought mechanisms to reduce the number of vehicles arriving at its plant each day. One proposal was for a carpooling system that would facilitate commuters sharing cars. This proposal lead to the development of the system discussed here.

Carpooling

Carpooling, (also known as ride sharing in the US) refers to a number of people making a journey in a single vehicle. This requires the identification of groups of people making similar journeys at similar times. Carpooling has long been used in the USA and various public initiatives exist to promote the use of multi-occupancy vehicles. A number of telephone based services exist to match potential drivers and passengers. The web has also been used to promote carpooling for example the services at www.rideshare.com (USA), www.carshare.ie (Ireland) and www.carpooling.com (Switzerland). Despite these initiatives, carpooling has not been very successful for social reasons, as most people find it more convenient to travel alone. Indeed there is some evidence that carpooling is actually declining. Against this background of public resistance to carpooling, it is important that any carpooling scheme can be operated with a minimum of inconvenience to those involved.

While some Intel employees share their cars, this number falls short of the potential. Consequently, some mechanism is needed to group people making similar journeys. A computer-based system was proposed which would conveniently match drivers and passengers. An important objective of this system was that it be totally automated, as personnel were not available to operate it. This would distinguish our proposed system from the web-based services described above, as these have limited interactivity and generally rely on human intervention in

the matching process. A further objective was that the system be completely general in nature, so that it could be used by other employers in Dublin or by Intel plants elsewhere.

Design Issues for the System

The various applications of information technology (IT) exploit three distinctive capabilities. The computational ability of IT, reflected in the name computer, allows the use of complex mathematical techniques. The storage capability, emphasized in database applications, allows ready access to large volume of data. Finally, the communications ability of IT, the basis of the current enthusiasm for Electronic Commerce, allows information flow to diverse locations. The carpooling system discussed in this paper exploits all of these capabilities. It allows a distributed population to communicate with the system, which stores their requirements in a database and then uses a modelling technique to come up with a good solution for the carpooling problem.

This type of system is related to a variety of E-Commerce systems that attempt to provide a link between different members of the public, for example auction sites. These attempt to minimize the transaction cost of an exchange between two users of the system. Carpooling is an example of an activity with considerable public benefit where those directly participating do not capture all of the gain to society. In this context there is a limited payoff to those involved so it is desirable that the system impose as little '*cost*' on participants as possible. This is reflected in our design for the system that imposes minimum operating requirements on its operator while being easy to use for potential car-poolers. Because of the broad public benefit, city authorities may be willing to support a system of this type. This support is more likely to be a grant towards its capital cost rather than the provision of continuing support for current costs.

The design of the system was influenced by a number of factors. The need for total automation is reflected in the direct user input of data over the web and the use of automatically generated emails to communicate the system output to the user. The need for a general design suggested the use of general-purpose software components rather than the use of custom code based around limited assumptions about the Dublin area or Intel's operations. These objectives were achieved by the use of the Transcad GIS to handle the geographic data in the problem. This allows a variety of map formats to be imported into the system. Standard database tables are used for the storage of other types of data, allowing easy interaction with other software. A web plugin is employed for the display of mapping data at the desktop; Autodesk Whip was used in the case of the prototype.

The system needed to capture data on the location of each user and communicate the solution to the users. This presented significant issues of privacy. A poorly designed system might provide a database of users' home addresses and what time they left for work at each day and might make this data accessible over the web. A number of design decisions attempted to reduce the potential privacy implications of the system. The system was designed to identify users using reference numbers derived from their employer. It is intended that email addresses or work telephone numbers are stored in the system rather than home telephone numbers. The precise home address of the user is not stored at all, when the user clicks on an onscreen map, they identify their approximate location in terms of the nearest significant location on the road network. This might mean identifying the entrance to a residential estate from a main road, but would not identify the precise house number at which the user lives. This information is sufficient for the automated matching and it is envisaged that the users themselves will reveal their address when making arrangements with the other people matched with their vehicle. These arrangements should ensure that that no information of value is available, even if a hacker was to break into the database.

The System

The system is composed of two loosely coupled sub-systems (Figure 1). The loose coupling between the components of the system is designed to allow it the possibility for the different components to be located on different servers on the Internet. The data capture system collects details of the user locations and travel times. These details are accumulated and the solver system is used to group passengers and drivers. This subdivision could assist in privacy as it could allow each employer to hold data on their own employees, passing only the minimum set of data to the solver system, located elsewhere on the Internet. This would assist in meeting the requirements of data protection legislation.

The data capture system provides a web-based application to allow potential users of the system to identify their needs. These come into two categories, the user identifies their location by clicking on an onscreen map, and the system then identifies the junction nearest the location identified and associates this location with the user. This data is incorporated in the relatively permanent user data, which also includes a reference number, user email address, and their place of work. The user also uses a web-based form to input more frequently changed data such as the times when they want to travel. Data is stored in database tables that can be accessed by the solver module.

The solver model is based on Transcad GIS software combined with a heuristic based algorithm for matching purposes. This reflects the need for GIS software to be augmented with specific routing algorithms. The Transcad GIS was used as it provides specific support for routing problems. Transcad contains a programming

Figure 1: Carpooling system

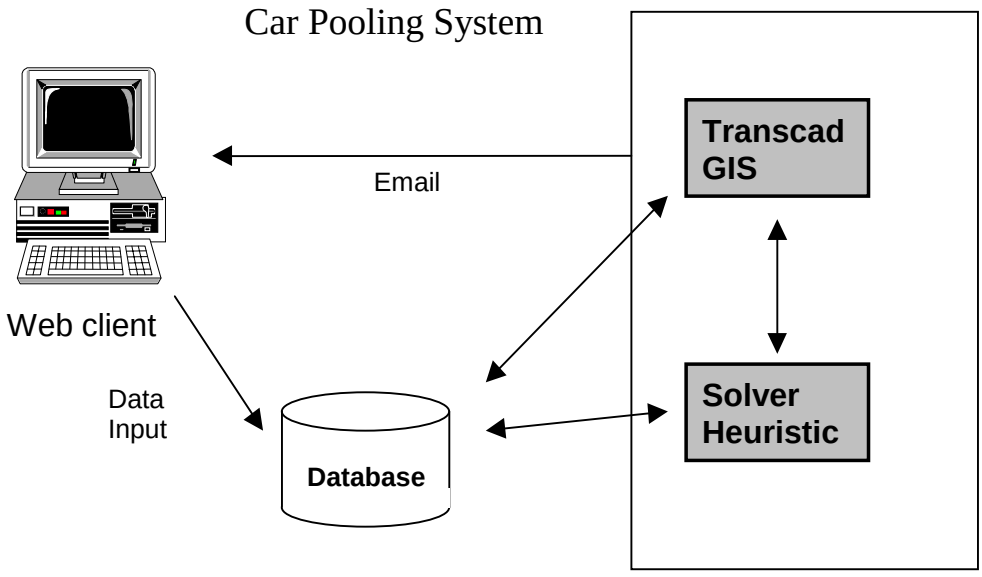

toolbox, the GIS Developers Kit, which allows easily development of customized applications. The solver program interacts with the GISDK to generate a time-based matrix of travel times between the driver locations, passenger locations, and travel destinations in the problem. The Transcad GIS allows the modelling of a variety of realistic road network conditions, for instance one way streets, no right turns (vehicles travel on the left in Ireland) and varying speeds per road segment. Where appropriate data is available, these capabilities allow realistic travel times be obtained for journeys.

The matching algorithm is heuristically based and attempts to minimize the marginal increase in the travel time imposed on drivers who have to pick up passengers. Each driver can set individual parameters for how much extra time they are willing to spend to pick up passengers. The extra time can be defined as a percentage of the time normally taken to travel to Intel or as an additional number of minutes. Each driver can also specify the number of passengers they are willing to carry, from one to three. Further restrictions can be added, for example to facilitate women who want to travel together or a smoker/non-smoker division. There are also a number of groupings of employees based on time, as some people work on a shift basis. Where you have a day shift and a night shift, the set of users is effectively divided in two for matching purposes. This reduces the size of the practical problem to be solved.

When a solution is derived, the system will automatically email the people involved. At this point the people concerned may or may not go ahead with carpooling in the groups suggested by the system. This presents a possible source of difficulty as the actual availability of vehicles may not correspond to that recorded by the system. With this in mind, a certain amount of redundancy is built into the operation of the system. For instance, emails are sent to both the driver and the passengers asking them to confirm that have gone ahead with the carpooling arrangements proposed by the system. In order to confirm the arrangements made, the user is asked to complete a web-based form. Email reminders can be sent repeatedly until a definite response is achieved.

Conclusion

At this point, the system has not gone into full operation. Tests have been conducted on the algorithm with a simulated distribution of user locations typical of the Intel workforce. Testing revealed that, depending of the parameters used, a sixty- percent reduction in the number of vehicle journeys could be achieved. Testing is taking place with students and users in Intel on the web-based map interface, to ensure maximum user-friendliness in the completed system. These tests suggest that the system will fully meet its technical requirements. Further research is required to indicate whether the system is sufficiently easy to use to increase significantly the proportion of people willing to carpool.