

IGT

Intelligent Grouping Transportation

Analysis and Background Information

Joe Olmi

www.taxibus.org.uk

joe@taxibus.org.uk

The Rationale for IGT

Throughout the world, road traffic congestion is becoming a critical problem. Congested roads dramatically increase journey times, making travel often stupidly inefficient. Being caught in slow moving traffic not only wastes time, but frequently puts drivers in aggressive moods. At the end of a car journey there is the added problem of parking: in many towns and cities finding a parking place can be very troublesome. Too many vehicles compete for limited parking space. The huge amount of time lost through traffic and parking delays costs the United Kingdom an estimated £20 billion each year. There are also enormous costs for maintaining and expanding the road networks to support the increasing levels of traffic. Traffic congestion significantly degrades the quality of life for everyone - a fact not so easily quantified, but one which is directly perceived by most people in their daily lives without the need for figures or statistics.

Pollution is a further and very pressing problem of high traffic levels. Engines that run on fossil fuels such as petrol or diesel emit many toxic pollutants: carbon monoxide, nitrogen dioxide, sulphur dioxide, benzene, formaldehyde, polycyclic hydrocarbons, lead, and small particle matter. Some of these pollutants react together in warm sunshine to form the poisonous gas ozone, the chief component of summer smog. In London, vehicle exhaust emissions are responsible for a staggering 90% of the air pollution.

Vehicle exhaust emissions are a substantial source of greenhouse gases. Fossil-fuel burning engines generate large quantities of carbon dioxide, a gas which though not toxic to human health, is the principal gas implicated in global climate change. The Transport Studies Unit at Oxford University estimate that at least 25% of Britain's carbon dioxide emissions originate from transportation, and in the US the figure is closer to 33%. Road traffic is thus very detrimental to global climate stability.

Traffic congestion, parking congestion, lost time and earnings, stressful journeys, road rage, degradation in the quality of life, noise pollution, air pollution, life threatening medical conditions, global warming: these are all problems of high traffic levels. How can traffic volumes be meaningfully reduced?

Problems of Existing Public Transport

Improvement of conventional public transport facilities is frequently claimed to be the best approach for reducing traffic levels. Public transport is unarguably a vital mode of travel, and road traffic levels would certainly soar without it. However public transport, in its existing forms, can never compete with the private car. The private car is unbeatable in terms of door-to-door convenience, all-weather comfort, personal security, flexibility of itinerary, ease of carrying goods or luggage, and so forth.

Compared to the private car, travelling by public transport frequently requires complex pre-planning (consulting timetables, examining fare options, checking

journey times, working out how to get to and from the train stations or bus stops). This planning becomes especially time consuming when travelling to new and unfamiliar destinations. For many people it will always seem easier to jump into the car. This fact is borne out by the Department for Transport's statistics: in the UK a massive 62% of all journeys take place by car, compared to only 6% by bus and a paltry 2% by train and underground railway. (Walking accounts for 26% of trips.)

These statistics further reveal that, for journeys of less than 25 miles, travelling by car is around twice as fast, on average, as the same journey by bus or train, when the overall door-to-door travel time is taken into account. Public travel takes so much longer than the car because of the necessary wait for the transport vehicle to arrive, and also because of the time consumed by trips to and from the station or bus stop at either end of the journey.

Given these facts, it is not surprising that the car is the first choice for transport. Clearly if public travel is to be made more attractive so as to compete with the private car, the sources of this wasted time must be eliminated, as must the complex pre-planning process, the inflexibility of itinerary, and the many other problems of public transport. Furthermore, as we saw, the private car supplies the bulk (62%) of the daily passenger journeys in the UK, so the car cannot be replaced until another mode of transport is introduced to provide for this amount of passenger journeys.

The Breakthrough Solution

IGT is the breakthrough solution to transport problems and traffic congestion. Its various modes of operation will radically reduce traffic levels, yet amazingly IGT requires no changes to road layouts, highway laws, existing public transport operation, or any transport infrastructures: it is a minimal-impact system, and one which easily coexists with other forms of transport and other traffic control schemes. This solution places absolutely no restrictions or additional costs on motorists. In fact in one mode of functioning it actually helps drivers recoup the costs of running their vehicles. Most importantly, IGT decreases the number of vehicles on the roads, yet it does not decrease the total number of passenger journeys taking place, thus maintaining the mobility of people.

IGT has two main configurations: taxibus transport, and car pooling transport. IGT can run either of these configurations separately, or can run them both at the same time.

The taxibus is a revolutionary and entirely new mode of public transport, perhaps destined to become the primary mode of travel in the 21st century. The taxibus is a road transport vehicle that generally conveys passengers from door-to-door, thus providing a transport service comparable to that of a taxi cab, yet one which may profitably operate with fare costs similar to those of a bus. The taxibus is the first mode of mass public transport that can not only equal the convenience of the private car, but as will become apparent, may even exceed it.

Analysis suggests that introducing a fleet of taxibus vehicles can massively reduce the quantity of traffic on the roads: a remarkable fourfold reduction in the number of cars on the roads is obtainable. The taxibus can swiftly curtail traffic congestion, air pollution and greenhouse gas emissions with unparalleled efficacy; for these reasons it is anticipated that taxibus transportation will be rapidly adopted in towns and cities around the globe.

The car pooling configuration of IGT, though not quite as revolutionary as the taxibus, nevertheless manages to provide a very effective means of matching prospective passengers with car pool drivers having compatible itineraries, and can provide a monetary incentive to encourage drivers to give rides. Although car pooling, even in this present incarnation, is an informal mode of public transport, it can certainly help further decrease the quantity of cars on the roads.

In both taxibus and car pooling, the transportation capabilities of IGT are completely scalable: IGT can operate in a village where it might provide just a few hundred passenger journeys each day, or in a major city where it could comfortably handle 10 million or more daily passenger journeys.

Note that IGT is highly practicable, both technologically and financially. For the most part it uses existing and established technologies, most of which are already in place or already in operation; as a consequence IGT can be implemented remarkably cheaply, and without changing any of the physical or technological infrastructures of a town or city. IGT is also cheap to run, and may quite easily be operated at a profit.

Basic IGT Concepts

The basic concepts of IGT are in fact extremely simple and can be understood purely in terms of passenger logistics. Computer technology is employed merely as a means to an end. IGT is actually based on two transport methodologies which we shall refer to as:

Adaptable Routing and Intelligent Grouping

Adaptable routing applies to any transport vehicles whose route is completely flexible and can be modified at any time to accommodate a traveller's particular journey requirements. Adaptively-routed vehicles provide a transport service that collects a traveller from his starting point address and conveys him to his destination address. This contrasts to fixed-route transport vehicles such as buses and trains which do not alter their routing to accommodate the traveller's itinerary. Ordinary taxi cabs use adaptable routing, as of course does the private car.

Intelligent grouping involves the placement of disparate travellers, who happen to have compatible itineraries, onto the same transport vehicle. Fixed-route vehicles such as buses and trains automatically use intelligent grouping; though in these cases, no great intelligence is involved, as the conveyance of

passengers with compatible itineraries is obviously intrinsic to fixed-route transport vehicles.

In the case of adaptively-routed transport vehicles, however, the intelligent grouping of travellers is by no means automatic and to achieve it requires precise orchestration of both travellers and transport vehicles. It is a complex undertaking to operate transport vehicles that simultaneously employ adaptable routing and intelligent grouping. Nevertheless, this is exactly what IGT does.

Adaptable routing is an important characteristic of IGT because it provides travellers with door-to-door travel from their starting point to their destination address, thus creating a convenient and highly attractive mode of public transport.

Intelligent grouping is an equally important characteristic of IGT as it enables transport vehicles to carry many disparate travellers at once, allowing passenger fare costs to be kept low as several travellers share the same transport vehicle.

The simultaneous use of adaptable routing and intelligent grouping in IGT is fundamental to reducing road traffic, as travellers who might have otherwise journeyed in their own individual cars can climb aboard an IGT transport vehicle and get exactly the same door-to-door service that their cars would have provided.

Precise Definition of Intelligent Grouping

DEFINITION: In an adaptively-routed transport vehicle, a set of travellers are said to be 'intelligently grouped' when the itinerary of any one traveller does not force the transport vehicle to significantly deviate from, and thus increase the journey times of, the itineraries of the other travellers aboard the vehicle. Obviously, a set of travellers can only be intelligently grouped when they have reasonably compatible itineraries.

IN OTHER WORDS: Any individual traveller would like to be conveyed on the quickest possible route from his embarkation to his destination - the route he would normally take in his own private car. The above definition of intelligent grouping states that in an adaptively-routed transport vehicle, a set of travellers are said to be intelligently grouped when each traveller's individual journey takes not much longer than this quickest possible route.

This is a qualitative definition of intelligent grouping, but we can be more precise and actually measure intelligent grouping quantitatively, as follows.

It can be shown that for each set of travellers intelligently grouped in an adaptively-routed transport vehicle, there is an **optimal transit route** which represents the quickest way to convey these travellers in accordance with their itineraries. For such a group of travellers, an important parameter is one called the **Compatibility Index** which quantitatively measures intelligent grouping. The Compatibility Index value for given any set of traveller itineraries is the

average increase in each traveller's journey time when these travellers are conveyed via this optimal transit route. This increase is in relation to the journey time that would result if each traveller were to travel on his itinerary via the most direct route, such as would normally be taken in a private car. Thus the Compatibility Index quantitatively measures the itinerary compatibility of a set of travellers.

The lower the value of the Compatibility Index, the more compatible the itineraries of the set of travellers. Perfect intelligent grouping has a Compatibility Index of 1. Compatibility Index values in the range of 1 to 1.3 represent excellent intelligent grouping, with values between 1.3 and 1.6 representing an acceptable level of intelligent grouping. A Compatibility Index approaching 2 or higher represents an increasingly inefficient and a generally unacceptable level of intelligent grouping.

Note that the Department for Transport statistics show that for trips of less than 25 miles, the total journey time when going by bus or train is a factor of 2 slower than the car, when the overall door-to-door travel time is taken into account.

By our definition, if we have a Compatibility Index of 1, this is as fast as travelling by car along the most direct route. Thus if the taxibus can operate with a Compatibility Index typically around the 1.3 mark, then it will come quite close to the speed of the private car and easily beats the bus or train (which correspond to an Index value of 2 by the Department for Transport statistics).

The Taxibus Fleet

The taxibus fleet could contain many types and sizes of transit vehicle. Vehicle size brings advantages and disadvantages. Larger transit vehicles can transport more people (and using just one driver, so it is also more economic on personnel). However, higher passenger numbers equate to longer journey times as a result of the extra pick-ups and drop-offs en route. By contrast, smaller vehicles transport less people, but do so more rapidly. Larger vehicles are efficient at transporting passengers that have identical embarkation or destination points, for example, passengers travelling from an airport to city centre hotels, or passengers travelling to a large public event such as a music concert or football match. Larger vehicles are also more appropriate for inter-city travel (if IGT is extended beyond just urban confines). Smaller transit vehicles are better at driving down narrow residential streets, picking up and dropping off passengers from their homes. It is thought that an urban taxibus fleet should comprise mainly smaller-sized vehicles, in the range of 8 to 12 passenger seats (in other words, the size of a multi-purpose vehicle or people carrier, up to the size of a minibus).

The journey range of each taxibus vehicle may vary. Some vehicles might be restricted for use within specified city areas; other vehicles might cover the whole of a city; still others might run from city-to-city, or even cover an entire country.

Occasionally a passenger's journey may be split into two legs, using two different taxibus vehicles. Typically, passengers travelling city-to-city by taxibus may have their journey thus split: a local taxibus vehicle first collects the passenger from his current address, and later the passenger is transferred to another (perhaps larger) taxibus vehicle for the inter-city portion of his journey. The IGT computer system will, however, always try to transport passengers on a single taxibus if possible.

An enhancement to further improve the efficiency of IGT would involve the computer system deploying its taxibus fleet according to the time of day. During commuter rush hours the computer system will tend to pull all its taxibus vehicles into operation; at quieter times such as during the night, many taxibuses will rest in their depots (small taxibuses might be parked outside the residence of the taxibus driver). In the morning rush hours, the computer system will ensure that plenty of taxibuses are mustered around the suburban regions ready to take passengers into the city centre. In the evening rush hours, the computer system will orchestrate the reverse. The IGT computer system will maintain a database of the statistically averaged passenger flux in all regions of the city, for all times of day, and for all days of the week. (The passenger flux from a given area is defined as the number of passengers per unit time requesting taxibus transport in that area). Using this passenger flux database, the IGT computer system will pre-emptively marshal its taxibus fleet so that the vehicles are fully prepared and positioned to meet temporal variations in travel demands (the computer system will pre-emptively re-position taxibuses that are currently empty). The computer system will also act on any notice that it is given of large public events, pre-emptively marshalling its fleet in order to handle the crowds.

Later in the evening and at night the IGT computer system will withdraw larger taxibus vehicles from the roads, keeping just the smaller taxibuses in operation. This will save fuel, keep noise levels down, and further decrease road congestion, as smaller vehicles take less space and cause less blockage when pulling over to pick up or drop off passengers. One of the problems with regular buses is that large vehicles are needed to deal with the peak hour passenger flux, but at other times these vehicles remain half empty. The taxibus fleet responds much more dynamically to variations in passenger flux.

Who would run the taxibus fleet? Many operators could be accommodated by franchising different organisations to run taxibuses. As long as each taxibus vehicle runs under the control of the IGT computer system, it does not matter who owns or drives the vehicle. Even an individual driver may purchase a taxibus vehicle and enter into a franchise; existing taxi cab or mini cab drivers might consider this. At the other end of the spectrum, large companies might operate a fleet comprising thousands of taxibus vehicles. IGT can accommodate all.

The IGT system may also accommodate specialist taxibuses: for example, large shopping centres or supermarkets might operate a small local fleet of taxibuses which would respond only to passengers headed for the shopping centre. Containing ample space for purchases, these taxibuses might even offer free rides to customers. Other types of specialist taxibuses could include ones designed for transporting particularly large or heavy items. There is even

the possibility of including first class taxibus vehicles in the fleet. First class vehicles would levy higher fares, but would contain larger and more comfortable seats, and provide additional services such as, for example, airline style back-of-seat Internet access terminals. A first class taxibus vehicle might even be designed to comprise several small compartments which are separated or semi-separated from each other, thus providing greater privacy for passengers carried. In general, the greater the diversity of taxibus vehicles operating, the better IGT is able to satisfy specific transport requirements.

Taxibus vehicles could advantageously be powered by hydrogen fuel cell engines. In certain configurations these engines are completely pollution-free: their exhaust product is actually pure water vapour. No pollutants or greenhouse gases are emitted. Hydrogen fuel cells create electricity which then drives an electric motor: vehicles with these engines thus run very silently. Large investments in developing hydrogen powered vehicles have been made, and they are fast becoming a viable technology. The main problem is the lack of fuelling stations. However, proponents for hydrogen fuel argue that at this initial stage, the rollout of fuel cell technology is ideally suited to public transport vehicles. Such vehicles are usually refuelled at their depots: this small number of depots can be equipped with hydrogen fuelling stations at minimal expenditure - thus providing an extremely cost effective way of cutting air pollution and greenhouse gas emissions.

IGT Car Pooling

IGT can also operate a very effective system of car pooling. Car pooling is an informal means of transportation which enlists the help of private vehicle drivers to furnish travellers with a door-to-door service similar to that supplied by the taxibus. It is hoped that the number of passenger journeys arising from car pooling will be a useful addition to the number of passenger journeys provided by the taxibus fleet, so that car pooling will further contribute to defeating traffic congestion.

IGT car pooling works as follows. To become a car pool driver, it is first necessary to set up a personal account with the IGT system. A user name and password are allocated to every registered user of the system. These are keyed into the car's dashboard display device, and subsequently transmitted to the IGT computer, in order to establish the identity of the driver to the IGT system. The driver then informs the IGT system of his intended destination and how many passenger seats are available in his vehicle. The IGT computer system always knows the car pool vehicle's current location through the GPS electronic positioning system in the car. Once the IGT computer system receives this information, the driver's vehicle is logged as available for car pooling, and the the driver may begin his journey.

Let us now turn our attention to the prospective car pool passenger. When a prospective passenger submits his journey request for a car pool trip, the IGT system will begin searching for a suitable car pool vehicle, seeking one that is already in transit, is located close to the passenger, and has an itinerary that can easily accommodate the passenger's desired destination. As soon as a

suitable car pool vehicle is found, its details are displayed on the passenger's cellular telephone. These details include: the car's distance from the passenger (in minutes), an estimate for the total car pool journey (in minutes), and the cost of this car pool journey. Should the IGT system find more than one suitable car pool vehicle, a list of vehicles will be detailed on the passenger's cellular telephone. The passenger then simply selects any vehicle from the list in order to proceed with the car pool trip. Note that the IGT system will only select car pool vehicles whose itineraries very closely match the passenger's, so that there is minimal inconvenience to the driver.

As soon as the passenger selects a car pool vehicle from the list, the IGT system will send a message to that vehicle's driver, requesting him to pick up and convey the passenger. The driver can accept or decline this journey request. If he does not respond to the request within a short time, the IGT computer system assumes that the driver has declined the request, and the passenger will be so informed. The passenger may then try another car on the list. Once the passenger finds a car pool driver who accepts his journey request, the intelligent grouping module will devise an optimal transit route which allows the car pool driver to efficiently pick up, convey and deliver this passenger. An in-car GPS street navigation system then directs the vehicle driver along this optimal transit route. By using in-car street navigation, the driver is guided to the waiting passenger with pinpoint precision.

A real-time countdown to the estimated car pool vehicle arrival time is shown on the passenger's cellular telephone. This is useful in conditions of bad weather, where the passenger can remain under shelter or indoors, coming out into the street just before the car arrives.

The principal incentive for a driver to participate in car pooling is monetary. For every car pool passenger carried, the driver receives a fee. For speed and simplicity, no money is exchanged in the vehicle: the IGT computer system handles all financial transactions automatically. Every person registered with the IGT computer system will have their own IGT monetary account, and after a car pool journey, the IGT system credits the driver's monetary account with his fee and debits the passenger's monetary account by an amount which represents the fare for his journey. However, many car pool drivers will not purely be interested in the financial reward; perhaps they also enjoy the company and conversation provided by passengers, or might just be curious as to whom they will meet. Should legislation be introduced that allows car pool drivers to use bus lanes (as do taxi cabs) or have access to free city-centre parking, or similar privileges, then the incentive to carry passengers is further reinforced.

The fare charge to the passenger and the fee received by the driver is calculated by the IGT computer system. There is a minimum set fare for a car pool journey (just as there is in taxi cabs) after which the cost is primarily based on the distance travelled. The controlling computer system will use the distance between the passenger's starting and destination point to calculate the fare. The total fare cost, however, will also include an amount related to the additional time added on to the driver's own journey as a result of diverting to collect and deliver the passenger: this ensures that the driver receives fair recompense when a larger route deviation is required to convey a passenger.

It is thought that car pool fares should be set at approximately two or three times the price of an equivalent bus ride, making car pooling a remarkably cheap form of door-to-door transport, yet from the driver's perspective, providing sufficient reward for carrying passengers. Note that when conveying more than one passenger, the fee received by the driver would be proportionally higher; a car driver can thus make a good profit from his journey by providing carriage to a series of passengers, or a group of passengers travelling together.

Intelligent Grouping Module

The intelligent grouping module (IGM) is the operational nucleus of the IGT software. When a new passenger submits his journey request, this module scans all available nearby transit vehicles in order to calculate which has an itinerary compatible to that of this new passenger.

There now follows an intuitive description of the desirable features the IGM might possess and how this IGM might handle certain specific situations.

First of all, it must be appreciated that the IGM will be performing an enormous amount of computational work during the normal operation of IGT. Consequently, it is important to ensure that the mathematical programming of the IGM's algorithm is as computationally efficient as possible. In particular, the IGM should be able to make rough estimates of itinerary compatibility with minimal computational effort. These rough estimates will enable the IGM to rapidly scan for good itinerary matches. Only when the module narrows down to a few good matches would a more accurate itinerary compatibility calculation be applied to these selected matches, in order to arrive at a precise Compatibility Index value for these few transit vehicles. Usually the vehicle with the very lowest Compatibility Index value will be selected.

The IGM always seeks to produce the best intelligent grouping possible, ideally finding a transit vehicle that can incorporate a traveller's itinerary without adding any extra time costs whatsoever to the itineraries of the other travellers aboard. This ideal solution will often be possible: as anyone who drives a car knows, between any two locations there are usually several alternative routes that are more or less equally time efficient. Thus when faced with a new traveller journey request, the IGM will first check if there are any transit vehicles that can accommodate the traveller simply by switching the vehicle to an alternative route, as such a transit vehicle could convey the new traveller without incurring any delay whatsoever to the journeys of existing on-board travellers.

This is the ideal case: new travellers are accommodated by an intelligent choice of route, and no time is wasted. However this perfect situation will not always present itself, and in general there will be a time cost for picking up and transporting new passengers. The IGM decides whether this extra time cost is acceptable or not, and this decision, though handled automatically, is carefully weighed up. A huge diversion with a large time cost to pick up just one single

passenger must be balanced by the inconvenience caused to the other travellers already on board the transit vehicle. The IGM will need to be carefully fine-tuned in order to get this balance right, so as to maximise overall transportation efficiency of the IGT fleet, but without inconveniencing any particular passenger.

The overall performance of the transit vehicle fleet is a key consideration: balancing the demands of routing efficiency against the need to pack as many travellers as possible into each transit vehicle. It must be understood that these two different demands can be in conflict, but may also be in harmony. The IGM aims for harmony: it tries to place travellers with highly compatible itineraries into the same transit vehicle, so that good routing speed and efficiency are attained, even though many travellers are simultaneously transported. Itinerary compatibility typically occurs when the embarkation and disembarkation points for all travellers are reasonably aligned on a direct path, thus satisfying both the need for routing efficiency and the need to convey a sufficient quantity of travellers. Though complete accord between these two requirements will not always be possible, the IGM always aims to achieve it.

When good accord between routing efficiency and passenger packing cannot be attained, the focus shifts to striking a balance between these two demands. This compromise will need to be carefully weighed up. It is important not to pack on too many travellers, as this may result in an excessively high value for the Compatibility Index of the transit vehicle, and a value greater than say 1.6 will not make these travellers very happy because their journeys will be quite elongated. If a new passenger would increase a transit vehicle's Compatibility Index to an unacceptably high value, then it may be better to place that passenger on another transit vehicle, if possible. On the other hand, it is important not to leave a passenger stranded: if there is only one transit vehicle currently in the vicinity of a new passenger, even though that transit vehicle would need to run a significant diversion to pick him up, it may be necessary to thus divert the vehicle in order not to leave this passenger waiting an unacceptably long time before the next suitable transit vehicle arrives. These sorts of situations will typically occur in the small hours of the night, or in smaller towns or rural areas, where there may not be so many transit vehicles available. During the day, and in major cities, there will be numerous transit vehicles, and the IGM will rarely need to be concerned with keeping a passenger waiting.

As well as the routing efficiency and passenger packing balance, there is another important balance that the IGM must always take into consideration: that between routing efficiency and vehicle proximity. The IGM aims to place new passengers in transit vehicles with the highest itinerary compatibility; however it also tries to find a transit vehicle that is in close proximity to the passenger, in order to have the transit vehicle arrive for pick up within three minutes of the passenger making the journey request. These two requirements may be in conflict. There may, for example, be a transit vehicle whose existing itinerary commitments can very easily accommodate the new passenger itinerary, but whose current location is six minutes away from the passenger; there may, however, be another transit vehicle whose itinerary commitments are not such a good match, but this vehicle is located just one minute away from the passenger. The IGM must make a balanced choice between keeping

the passenger waiting six minutes for the more compatible transit vehicle, or placing him on board a transit vehicle within one minute, and accepting the compromise in routing efficiency (and therefore journey time). In fact, this balanced choice is informed by the prospective passenger's own preferences. If the passenger prioritises **speed of response** when submitting his travel request, then usually the nearest transit vehicle will be allocated to him; on the other hand, if the passenger prioritises **speed of journey** when making his travel request, then the vehicle with the best itinerary match will be selected to convey this passenger.

Another automatic yet balanced decision the IGM must make is whether to convey a passenger on a single transit vehicle, or split a passenger's journey using two or more different transit vehicles.

There will be many other balances that the IGM automatically takes into account. These various balances may have to be fine-tuned on a trial and error basis: adjusting a little, and observing the effect it has on overall transit vehicle fleet performance (this is where a computer simulation of the transit vehicle fleet becomes very useful).

Different routing techniques could be experimented with to see which best expedites the overall transportation process. For example, it is conceivable that the technique of location clustering might speed up certain journeys. Location clustering is the bunching of pick-up points and/or drop-off points within the same small area. This might work well for suburban commuting, where a transit vehicle will pick up passengers from a cluster of homes in a particular suburban region, run all the way into the centre of town on a main road without stopping, and deliver these passengers directly to their various places of work, these places again clustered in the same locale.

Another technique relates to travellers that have identical pick-up or drop-off points: if these travellers also have compatible itineraries, placing them on the same transit vehicle can further increase transit vehicle efficiency, because the vehicle will have less individual pick-up or drop-off points on its route.

For the best results in calculating optimal transit routes, the IGM must take into account not only the distance along a route, but also the speed at which a transit vehicle can traverse the various roads that comprise the route. Although such a **traffic-speed travel-time metric** (see below for definition) is more complicated to calculate, it yields many advantages. The complications come from the need to estimate the speed at which a transit vehicle can travel on each road.

A simple approach to estimating these speeds assumes that the transit vehicle will travel at the speed limit of each road traversed. Another simple approach - which would work quite well in traffic congested cities - assumes that on all roads, the transit vehicle will travel at the average traffic speed for that time of day (for example average traffic speeds in suburban London are around 20 mph for most of the day). A more sophisticated approach assumes the transit vehicle travels at the average traffic speed of each particular road, for the particular time of day, these average traffic speeds being calculated from previously-recorded electronic positioning data received from transit vehicles

in the fleet.

The most refined approach would use real-time electronic positioning data received from the transit vehicles fleet to provide an up-to-the-moment analysis of the current traffic speeds on all roads. This last approach is the one recommended, not only because it provides the most accurate estimate of traffic speeds, but more importantly, because it allows the IGM to automatically take any traffic jams into account when devising optimal transit routes: if the current estimated traffic speed on a particular road is very slow, then the IGM would, by the definition of the optimal transit route, tend to avoid routing transit vehicles via that road until the traffic situation improves. This means that the transit vehicles of IGT will tend to pre-emptively avoid traffic jams.

Note that once the IGM has grouped travellers in a transit vehicle and has devised a optimal transit route for that vehicle, the task of directing the transit vehicle driver along that route is delegated to the electronic street navigation module, which operates in a similar fashion to in-car satellite navigation systems that are familiar to many drivers. With the electronic street navigation module in control, should the vehicle driver want to change course slightly (in order to avoid a small traffic jam for example), or should the driver simply make a routing mistake, the electronic street navigation module will adapt to, and continue from, his new position and circumstances (such adaptability and flexibility is one of the useful features of in-car satellite navigation systems, as any driver that has used one knows).

However should there be, for whatever reason, a large change in circumstances such that the transit vehicle gets significantly displaced from the original optimal transit route, then, with respect to this transit vehicle's displaced present position, the original optimal transit route may no longer be the most efficient way of transporting the transit vehicle's passengers. So rather than letting the electronic street navigation module try to adapt to the highly displaced position, instead, the IGM will momentarily step in to re-optimize the transit vehicle's route. After the IGM has devised a new optimal transit route for the vehicle, based on its displaced position, the electronic street navigation module will proceed as normal, directing the transit vehicle along that new optimal transit route.

This route re-optimisation would be set to kick-in automatically whenever a transit vehicle becomes significantly displaced from its intended route. Thus, provided a transit vehicle keeps close to the optimal transit route originally devised by the IGM, transit vehicle navigation will remain under control of the electronic street navigation module; but should a transit vehicle, for whatever reason, significantly deviate from this optimal transit route, this will trigger the IGM to step in and re-optimize that transit vehicle's route.

Such re-optimisation may mean, for example, that the transit vehicle will be given a new path to follow, and may also mean that it will be instructed to pick up and drop off its passengers in a different sequence. Re-optimisation might also entail that one or more passenger pick-ups planned for that transit vehicle are now cancelled, with those waiting passengers now collected by another transit vehicle, the IGM having calculated that, under the present circumstances, this is the most efficient way to convey them.

Re-optimisation would also be set to trigger in other significant circumstances, such as when a transit vehicle gets delayed for a long time due to heavy traffic, or when the traffic congestion conditions have significantly deteriorated ahead along the current optimal transit route of the transit vehicle, now making that route less than optimal.

This complex juggling of circumstances is performed automatically in order to maximise the overall speed and efficiency of the transit vehicle fleet; transit vehicle drivers themselves are not involved (or aware of) this process. It is the IGT computer system that sweats; the transit vehicle drivers merely follow the resulting navigational instructions shown on their dashboard display device.

In summary: this section has considered some desirable features that the IGM might have. However this is not an exhaustive survey, and further performance-enhancing features may be added. The beauty of IGT is that when performance improvements are made to the IGM, these instantly alter the operation of the whole transit vehicle fleet. It is anticipated that the mathematical operation of the IGM will be perfected over a period of many years. Each city in the world that implements IGT can employ mathematicians and software engineers to try to further increase the transportation efficiency of the IGM: the great virtue of IGT is that any new performance-enhancing features devised for one city can be easily implemented in other cities, just as a software upgrade.

Real-Time Traffic Speed Module

Once a sufficient quantity of IGT transit vehicles are in operation in a city or similar region, it will be possible to assess the current average vehicle speeds on all roads in that city. This is achieved by a **real-time traffic speed module**, which is a software module that runs on the IGT computer system and examines incoming GPS electronic positioning data from all transit vehicles in order to estimate current traffic speeds. This module has the ability to automatically detect any roads with traffic jams or traffic blockages simply because the transit vehicle speeds, as calculated from their electronic positioning data, would obviously be very slow on these roads.

The real-time traffic speed module operates using a digitised street layout map of the city. Each time a transit vehicle in the IGT fleet traverses a stretch of road, the real-time traffic speed module will use incoming electronic positioning data from that vehicle to plot the vehicle's progress. By noting the time the transit vehicle takes to cover a known length of road, the module can determine the transit vehicle's speed on that road. Working on the assumption that transit vehicles travel at the prevailing traffic speed (when not picking up or dropping off passengers), the real-time traffic speed module can estimate the real-time current traffic speed on each road in the city.

When a large fleet of transit vehicles are operating in a city, one would expect several transit vehicles to traverse a typical city street in each ten-minute period. For every stretch of road, the real-time traffic speed module averages the transit vehicle speed values it receives over this time period to provide a

constantly updated road-by-road city-wide map of the real-time current traffic speeds.

For roads on which transit vehicles run only infrequently, the real-time traffic speed module would use other methods to estimate prevailing traffic speeds (for example, if a road only infrequently has transit vehicles running along it, then one might reasonably assume that this road has very little traffic in general, meaning that any vehicle travelling along it will be able to drive at the speed limit of the road).

With this real-time traffic speed data at its disposal, the IGT computer system will be able to automatically and pre-emptively re-route its transit vehicles away from blocked or heavily congested roads, thus cleverly avoiding delays (and also relieving the traffic burden at these already-congested points, which benefits everybody).

Testing IGT

The IGT taxibus is an entirely new form of transport, and must initially be tested on a small scale, firstly to examine its overall functionality, and secondly to optimise certain design and operational parameters.

Implementing the car pooling mode of IGT is a simple and inexpensive way of testing its general functioning: no vehicles need to be purchased, so the car pooling mode can be implemented at a very low cost. Once car pooling is up and running, it will yield valuable data on how IGT operates in reality. However, there are many differences between car pooling and the taxibus, and the taxibus transportation system needs to be tested separately.

Perhaps the cheapest way of testing the taxibus is by means of computer simulation of a city's travel networks. The software for such a simulation could be written without much difficulty, and would be extremely useful for observing the effects of adjusting various parameters, such as the number of taxibuses in the fleet, the passenger capacity of each taxibus, and the operational characteristics of the intelligent grouping module. By means of this simulation, optimum parameter values can be determined before setting up a real taxibus service.

Subsequent to computer simulation tests of the taxibus, the next step would be to implement a taxibus service on a small scale. This might best be accomplished in a city of less than half a million inhabitants, wherein a relatively modest fleet of taxibus vehicles can cover the whole town. It is very important to implement a saturation level of taxibus vehicles in any test. Saturation level is defined as the minimum taxibus vehicle density that can consistently provide a three minute response time to journey requests. If the taxibus density is below the saturation level, then the average response time to passenger journey requests becomes longer. Slow response times will put people off using the taxibus: nobody is going to find the taxibus quick and efficient if it takes 20 minutes for the vehicle to arrive. Thus IGT must be tested at saturation capacity in order to properly examine its viability as a public

transport service.

A small scale test is even possible in large cities. This is achieved by providing a saturation level, 24 hour taxibus fleet to a particular suburb in the city. Just a few hundred taxibus vehicles will be needed for saturation coverage. Travel on these taxibuses would be restricted to journeys within the selected suburban area, and journeys from the suburban area to the city centre and back. These limited transport options will not be able to accommodate all the travel demands of these people, but should encompass a good percentage of their normal journeys - enough to make the test meaningful.

Small scale testing in a large city may also be achieved by confining the taxibuses to the central area. The central London region, for example, bounded by the Inner Ring Road covers just eight square miles. A fleet of 200 vehicles constrained to this zone would yield a density of 25 taxibuses per square mile, which should be sufficient.

Note that the vehicles used for testing need not be specially designed or purchased. Large cars and minibuses can be employed. For testing purposes, the only essential modifications required are the fitting of a small computer, cellular network link and GPS device, and the displaying of a vehicle identification number (or name) on the outside of the taxibus. However, if the budget is available, a fleet of new, purpose-built vehicles will obviously make a better impression on the public.

Quai Door-to-Door Taxibus

Traditional public transport generally has its access points on main streets in the form of train stations and bus stops. However the IGT taxibus is a unique form of mass public transport which distributes its passenger pick-up and drop-off points more widely, including smaller roads as well as main streets. It is believed that this strategy will improve overall transportation efficiency - even if some time is lost during taxibus excursions into the back-streets.

Here we present some variations on this strategy that may help further streamline taxibus travel.

Since the taxibus conveys passengers on a door-to-door basis, a certain portion of its journey will be spent in excursions into residential or other back-street areas in order to pick up or drop off passengers. The intelligent grouping module always tries to streamline these excursions so that any diversion into the back-streets to collect or deliver a passenger will also act as a short-cut or cut-through on the overall taxibus route. However, not all back-street excursions can be made to double up as short cuts, and sometimes they will slow the taxibus journey down.

One way to eliminate this slow down is by confining all taxibuses to the main roads and larger streets only, with passenger pick-ups and drop-offs taking place exclusively on these main roads (just like with regular buses). This confinement to larger streets would be controlled by the intelligent grouping

module which would be set up to create optimal transit routes that avoided the back-streets. Although it no longer offers a true door-to-door service, the quasi door-to-door taxibus is still much more convenient and flexible than a regular bus service because it continues to allow complete route customisation according to each passenger's individual itinerary.

Using this quasi door-to-door approach, each main street would be provided with several taxibus stops containing a kiosk device into which passengers can enter their journey request, and at which they will wait for their requested taxibus to arrive. The disadvantage of the quasi door-to-door taxibus is that passengers must walk to the nearest taxibus stop in order to request and board a taxibus and, on disembarkation, passengers will typically need to walk the last few hundred metres to their destination. Thus this service is obviously not as convenient as ordering and waiting for a taxibus within the comfort of your own home, office, or suchlike.

One clear advantage of the quasi door-to-door taxibus, however, is that with passengers congregated in the street at taxibus stops, the taxibus will perform less individual passenger pick-ups and drop offs, as several passengers will tend to board and alight together. This will speed up the taxibus since there will be a reduced number of passenger stops en route. Furthermore, by picking up passengers in larger batches, it becomes feasible (and necessary) to use larger capacity taxibus vehicles, thus further increasing the fleet efficiency and at the same time reducing the total number of taxibus drivers needed to run the fleet.

A hybrid between the quasi taxibus and the true taxibus service is also conceptually possible - and is an excellent compromise between efficiency of the former and the convenience of the latter. In this hybrid system, the taxibus will pick up passengers directly from their current location, but will drop off passengers only on main streets or larger streets. Thus passengers may need to walk the final few hundred metres to their destination. This sensible compromise will be especially useful near one-way systems and no-through roads, where a taxibus might otherwise have to make a large diversion in order to deliver a single passenger. As an alternative approach, this hybrid system might be set to operate only at such awkward destinations as no-through roads and one-way systems, but otherwise deliver passengers to the door. There are many other variations on this quasi door-to-door theme which may also prove expedient.

Taxi Cabs and IGT

Regular taxi cabs are readily included in IGT, and there are two ways a taxi driver can benefit. In the first, the taxi driver uses IGT just to acquire customers, in a manner similar to the existing 'radio taxi' passenger procurement system. In this mode of operation, when a passenger makes a journey request, his cellular telephone, in addition to detailing the currently available taxibus and car pool travel options in the passenger's vicinity, will also detail the taxi cab options available. Should the passenger select a taxi cab option, the IGT computer system will respond accordingly: the journey

request will be detailed on the selected taxi driver's dashboard display device, and if the taxi driver accepts the journey request, the IGT computer system will guide him to the passenger, just as it does with car pool drivers. This taxi mode of operation is not strictly speaking an application of IGT since it does not follow the principle of intelligent grouping; nevertheless this taxi mode can easily be provided by the controlling computer system, and will be helpful for both taxi drivers and their passengers.

The second way a taxi driver can benefit from IGT is to operate his vehicle as a taxibus. In this taxibus mode of operation, the passenger fares charged will be lower - being priced similar to car pool fares; however the taxi driver will be conveying perhaps three or four disparate passengers at once, so their concurrently metered fares will add up to an amount which is comparable to a normal taxi fare. It is entirely up to the taxi driver how he wants to operate: he can function in taxi cab mode or in taxibus mode. The mode of operation can be switched at any time; the taxi driver would select the mode in order to maximise his business and profit.

Even if taxi drivers prefer to remain as regular cabs, operating only in taxi cab mode, it is anticipated that IGT will deliver taxi drivers more custom. Many prospective passengers using IGT will want to take the first and fastest transport option available: this will always be the taxi cab. Why? Because with taxibus or car pool travel there are two requirements to satisfy: the transit vehicle must be in proximity to the passenger and the vehicle must have itinerary commitments compatible to the passenger's itinerary. With the taxi cab there is only one constraint to satisfy: proximity. Thus a taxi cab vehicle will generally be the first available, and because it does not pick up other passengers en route, it will generally be the fastest available. For these reasons it is believed that IGT will deliver more business to taxi cabs. Whether the taxi driver's fee is paid automatically into his system-administered monetary account, or paid in cash within the cab, is a question that needs further consideration.

One advantage IGT offers the taxi driver is the ability to accept or decline incoming journey requests on the basis of the passenger's specified itinerary - this itinerary will always be detailed on the taxi driver's dashboard display device. This is useful when the taxi driver is finishing his working day and is about to return home: in these circumstances he can select a passenger with a destination close to his home, and thus conveniently include one more fare before he ends his shift.