Monitoring the Movement of Pedestrians Using Low-cost Infrared Detectors: Initial Findings

Jon Kerridge, Alistair Armitage, David Binnie and Lucy Lei {initial.surname@napier.ac.uk}

Napier University, Edinburgh, Scotland, EH10 5DT

Neil Sumpter {neil.sumpter@irisys.co.uk}

IRISYS Ltd, Towcester, Northants, UK, NN12 6AD

ABSTRACT

Monitoring the movement of pedestrians in everyday environments is difficult, especially if accurate data is required. Typically data pertaining to the origins and destinations of pedestrians as they move around a space can only be obtained either by locating many observers in the area under surveillance or by analyzing CCTV footage. The former is error prone and perturbs the space being analyzed; the latter is expensive in terms of the technology and time required. Technical advances in low-cost infrared detectors provide an opportunity to unobtrusively observe pedestrian spaces and determine individual pedestrian trajectories automatically. We describe the underpinning detector technology and show how arrays of such detectors can be used to monitor larger spaces. An outline of the algorithm used to create complete trajectories as pedestrians move between detectors is presented. A series of experiments is described where pedestrians were asked to move in a set of defined patterns in a controlled environment. Initial results from these experiments are discussed where we found that at Fruin Levels of Service A to C 93% of pedestrian trajectories could be tracked, which dropped to 79% at Level of Service D.

INTRODUCTION

Many governments and transport authorities are firmly committed to the development of an integrated transport policy. The success of this policy will depend on ensuring that the different modes of transport work effectively and efficiently together and this, in turn, will require the optimization of flows between modes. In most cases flow between modes is in the form of a pedestrian sector in a journey. Recent attempts to improve pedestrian flow have relied on monitoring methods and systems that are expensive to implement and limited in functionality. The utility and application of these attempts to date has thus been constrained by technological and financial requirements. There are many situations where knowledge of the number and movement patterns of pedestrians would be beneficial including:

- people in bus and taxi queues to ensure adequate vehicle deployment
- pedestrian density within rail station concourses or on platforms to prevent overcrowding
- pedestrian flow rates e.g. out of stations to enable the road vehicle management systems to optimise pedestrian and vehicular traffic movements
- passenger flow rates within airports to optimize the number of passport control staff available
- town planning authorities require pedestrian flow measures to monitor the viability and vitality of parts of a town.

It is the last of these that provides the motivation for the research and development that has been undertaken. The viability of an area is a measure of the number of retail outlets of different types that can be sustained by the local population. The vitality is a measure of how many people use those outlets. Obtaining this last measure is a specific objective of the research programme.

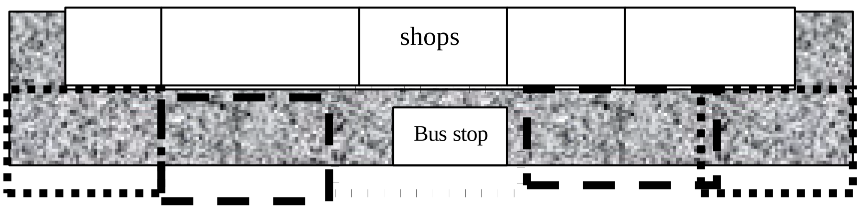

Figure 1. A Row of Shops to be Monitored

Figure 1 represents a retail area in which there are a number of single entrance retail outlets together with a bus stop. It also shows the area in five sections that can be monitored using the detector technology we have used. The goal of the monitoring is to determine the number of people who enter each shop and to determine the routes taken by people as they move through the area. This includes, for example, how many people leave one shop and then pass to another particular shop, or move to the bus stop and how many people pass the shops without entering any of them. Such observations need to be undertaken over an extended period if meaningful analysis is to be undertaken.

The goal was to produce a system that can be applied, over the course of a day, to:

- count individual people at various datum points within the space.
- determine the number of people alighting from and embarking on a bus.
- count the number of people entering and leaving each shop.
- make a reasonable attempt to follow a person as they leave a shop and enter another one to determine the number and type of intermediate goals in a journey.

In order to satisfy the goal the following research questions had to be answered:

- Can we easily track an individual from the field of view of one detector to the next?
- Can we dynamically create datum points over which we can collect various statistics?

PREVIOUS WORK

Many groups [19,20,21,23] have built image-processing systems, which extract pedestrian activity information from visible spectrum images. This technique relies on extensive use of large computing resources due to the large amount of information, which needs to be processed prior to the extraction of pedestrian data. In particular, Reading [19] has shown the difficulty of detecting people at Puffin crossings using visible spectrum cameras. Many workers have attempted to follow individuals from the field of view of one camera to another and this has proved to be a very challenging problem. One of the greatest assets of an infra-red device of the type used in this application is its ability to locate people without the host of extraneous data that plague visible CCTV images.

Willis [23] describes a method of extracting pedestrian trajectories from CCTV footage using a commercially available motion analysis tool that enables frame-by-frame determination of position and hence allows determination of specific measurements in the environment. Similarly, many researchers have attempted to capture pedestrian movements in a variety of different scenarios [2-10,13,14,16-18]. In all these studies, observation of the pedestrians involves either people watching the scene, or subsequent human analysis of videos of street scenes [9,10]. In either case this is a time consuming process and susceptible to error as the human observer is not capable of maintaining the degree of concentration required to eliminate all errors.

Previous studies have attempted to categorize pedestrian behavior into one of a few classes and have generated a simple density flow metric. There has been no study of detailed information on pedestrian flow due the difficulty in gathering accurate movement data over an extended period. It is very difficult for human of observers to track an individual to determine the shops they visited along the section of street. Hoogendoorn et al[12] describe a series of laboratory experiments to discover how pedestrians move around bottlenecks in the pedestrian environment. Each of the experimental subjects wore a different colored hat and adopts a previously defined behavior associated with the color of the hat. The experiments took place in a large hall with a video camera placed directly above the experimental area. They used standard image processing techniques to obtain trajectory data. In the work reported in this paper we are trying to obtain trajectory information from natural situations where pedestrians are unaware that they are being observed. The detectors we use successfully track individual and group flow and can work throughout the day and night and over a period of days or weeks. The results reported here were obtained during the course of a two-and-a-half year project.

TECHNOLOGICAL UNDERPINNING

The Irisys detector [11,15,22] is primarily designed for counting people moving back and forth across a datum line specified by the user. In this mode a single detector can be used, for instance, to count movements along a corridor or through a doorway. Multiple detectors can be aligned to create a larger counting line. In this mode detectors have been used successfully; to count pedestrian movements in and out of supermarkets, in station concourses, and over the Millennium footbridge in London during trials after the bridge was modified. A pair of detectors has been mounted on either side of an automatic door to monitor the movement of people in its vicinity. The aim is to determine the trajectory of pedestrians and only open the doors if a pedestrian is directly moving towards the door. In this mode the detectors are also able to determine if a pedestrian has actually moved through the doors hence holding the doors open for as long as is necessary. In addition to ensuring that the doors do not trap pedestrians, it also ensures that any climate control inside the building is not expensively wasted.

We have been able to extract more information from the detectors, in particular to record pedestrian trajectories [1]. In normal counting mode, the detectors can be used stand-alone, with total people counts being read out periodically, or connected to a simple counter. In trajectory mode, the detectors are connected to a computer that gathers the extra data and also timestamps the data from each detector.

The detector uses a 16 x 16 array of pyroelectric ceramic detectors to measure changes in temperature. Most infrared imaging systems measure absolute temperature, or use an internal chopper to measure the temperature difference between the scene and a known reference; this adds complexity to a pyroelectric detector. The detectors we have been using only measure temperature differences, and rely on the pedestrian being at a different temperature from the background. This has the advantage that the background disappears from the image, leaving pedestrians as clear targets. In Figure 2 the white areas indicate a person, with the darker areas indicating the background.

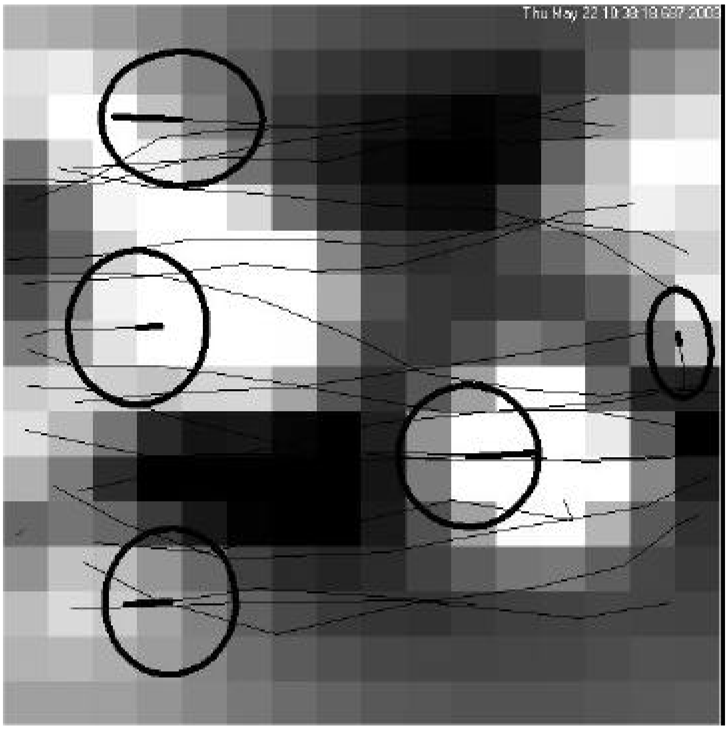

Figure 2. An image of a typical scene containing five pedestrians with superimposed trajectories.

The normal mode of using the detector is to mount it at a height of three metres or so above the region of interest vertically above the space to be observed. At this height, the detector covers a ground area three metres by three metres. The detector has an internal processor that undertakes image pre-processing. Likely targets have an ellipse fitted to them, and the center of the ellipse is then calculated and communicated down a communications link to a data-gathering computer. The X and Y positions are recorded approximately three times per second. These values are presented as a floating-point number in the range 0 – 16 for each axis. Each unit, being the width of a pixel, corresponds to about 20 centimeters on the ground. However, as the centering process effectively averages data from a number of pixels, the position is actually recorded with sub-pixel accuracy. **Figure 2** shows the ellipses that have been matched and the trajectories of these targets and other previous targets. The line within the ellipse shows the direction of travel. **Figure 2** also demonstrates that it is impossible to identify any individual from the image. The ellipse appears behind the white area indicating the person due to a time lag between the image data and the processed target information. Each person image comprises a hot part (white) followed by a (black) cold wake. The size of the white part does reflect the rotundity of the person. The grey pixels indicate moving arms and legs. It is not possible to distinguish anything about a person apart from their general size; for example, children cannot be distinguished from small adults.

There are several limitations in the detectors we have been using so far. The frame rate is approximately three frames per second. For many pedestrians, this corresponds to one reading every half-metre or so. Ideally, we would like a finer temporal resolution for some applications. However, a more serious limitation is caused by the difficulty the internal processor has in acquiring targets at the edge of an image. This can be seen in **Figure 3**, where a number of spurious artifacts appear, as if pedestrians had suddenly changed course at the edge of the field of view. Many of these can be eliminated using the status information provided by the detector system. However, this means that some readings at the edges of fields are dropped. The rightmost target is typical of such an artifact.

The detectors we used have now been replaced by an improved version, which overcomes the problems discussed previously. Details of the latest Irisys detectors are given on the 'Products' section of the company web site [25], with further background being given in the 'Technology' section. Several models of people counter are made, but the main differences are in the housing; two models are available depending on whether they are to be used indoors or outdoors. The results reported in this study used the earlier model of detector and we would therefore expect an improvement in the performance of the pedestrian tracking capability once we use the latest technology.

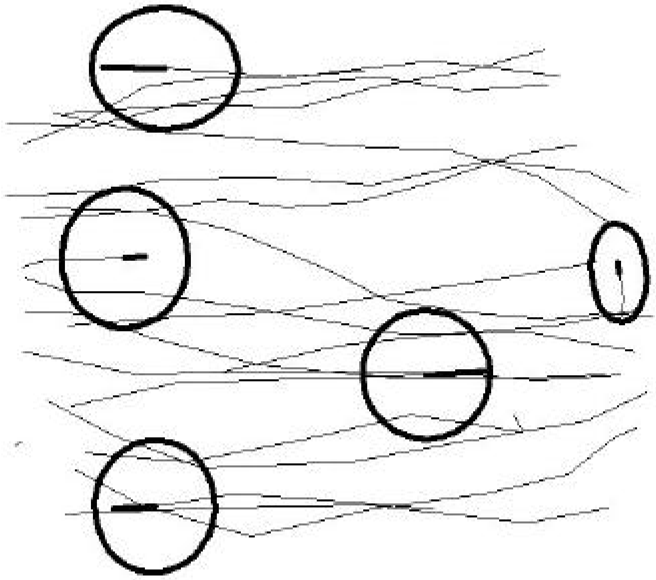

Figure 3. Highlighted Targets and Trajectories from Figure 2

MULTIPLE DETECTOR SYSTEMS

The main problem with multiple detector systems is tracking the trajectory from one field of view to the next. The nature of the infrared detectors makes it difficult to line them up precisely. We cannot guarantee the relative alignment of the detectors and thus the data analysis systems have to be able to overcome such a situation and cannot rely on a perfect alignment of the fields of view of adjacent detectors to a predefined specification. For each detector a set of data is captured in the central computer of the form shown in **Figure 4,** which represents a single person moving through the field of view of a detector. A record starting with C gives a time stamp for the data reading with an accuracy of milliseconds. A record starting with T gives the coordinates that have been fitted to a person by the detector. In this example a person starts from a position of [0.416, 4.789] at time Thu May 22 10:36:41.531 2003and leaves the field of view at location [13.064, 3.179] at time Thu May 22

10:36:43.000 2003, if we ignore data points with a status bit of 1. If there is more than one person in the field of view then for each timestamp (Type C) record there will be more than one target data point (Type T) record. If no targets have been detected a Type C, timestamp, record will be immediately followed by another and the detector will re-initialize the identification number back to the start of the series.

Timestamp		Target Data				
		Id	S	X	Y	Mode
C Thu May 22 10:36:40.906 2003						
C Thu May 22 10:36:41.218 2003	T	2	1	0.304	5.089	257
C Thu May 22 10:36:41.531 2003	T	2	0	0.416	4.789	261
C Thu May 22 10:36:41.859 2003	T	2	0	0.776	3.303	261
C Thu May 22 10:36:42.171 2003	T	2	0	4.256	2.860	261
C Thu May 22 10:36:42.375 2003	T	2	0	7.636	3.020	261
C Thu May 22 10:36:42.687 2003	T	2	0	10.831	2.749	261
C Thu May 22 10:36:43.000 2003	T	2	0	13.064	3.179	261

Id – Target Identifier; S – Status, 1 – invalid, 0 – valid
X – X-pixel co-ordinate; Y – Y-pixel co-ordinate
Mode – Values greater than 258 give additional data when S = 0
Values less than 258 confirm the nature of invalidity when S = 1

Figure 4 Data Record for a single person in the Field of View

Figure 5 shows the data coming from one detector in the middle of a data collection exercise. The data has been annotated to provide a commentary of the detected behaviour. Each detector collects its data independently and transmits its data stream to the central computer where the streams are collected in separate files. The task of finding out what happened is thus one which has to be undertaken retrospectively by merging the multiple data sets so we can track a person as they move from one detector to the next. Extracting the trajectory with a single detector is relatively easy as an individual is assigned an identification number. Thus we can determine the starting and ending time of an individual, their entry and exit position simply by analysing the data points for the identification number chosen. Thus in **Figure 5** individual 8 enters at 10:37:19.640 and exits at 10:37:21.156 where the last number represents milliseconds. During this time the individual moved along the trajectory [0.348, 4.119], [1.870, 3.540], [4.905, 3.435], [8.269, 3.525], [11.686, 3.554], [14.527, 3.613]. Movement in the y direction is small whereas, movement in the x direction is large as the individual moved across the field of view. If we presume that the person moved in a straight line then we can determine the average speed. The data points occur at approximately 300 ms intervals. If we assume 20 cm per pixel then:

the distance covered is (14.527 - 0.348) * 0.20 = 2.84 m

in a time of (21.156 – 19.640) = 1.36seconds

giving an average speed of 2.09 m/sec

The internal processor in the detector has a number of behaviors that have to be understood if the data is to be manipulated correctly. The identification number is unique to a detector. The sequence of identification numbers is reused from the initial value once the detector has been unable to detect any targets in the field of view for some time. This means that the lower identification numbers are reused many times during an observation, hence we have to undertake additional processing between the streams of data if we are to track people across multiple detectors. This is discussed in the next section.

	Id	S	X	Y	Mode	Annotation
C Thu May 22 10:37:19.640 2003						
T	7	0	11.985	6.104	261	*target 7 is moving across the field of view*
T	8	0	0.348	4.119	261	*target 8 has just started its transit*
T	9	1	0.009	8.926	257	*target 9 has just entered with an invalid status*
C Thu May 22 10:37:19.953 2003						
T	7	0	14.094	6.477	261	*target 7 continues*
T	8	0	1.870	3.540	261	*target 8 continues*
T	9	0	0.658	8.666	261	*target 9 gets its first valid co-ordinates*
C Thu May 22 10:37:20.296 2003						
T	8	0	4.905	3.435	261	*target 8 continues, target 7 has left*
T	9	0	3.188	8.534	261	*target 9 continues*
C Thu May 22 10:37:20.515 2003						
T	8	0	8.269	3.525	261	*target 8 continues*
T	9	0	6.134	8.532	261	*target 9 continues*
C Thu May 22 10:37:20.828 2003						
T	8	0	11.686	3.554	261	*target 8 continues*
T	9	0	9.438	8.472	261	*target 9 continues*
T	10	0	0.647	6.786	261	*target 10 has entered with a valid status*
C Thu May 22 10:37:21.156 2003						
T	8	0	14.527	3.613	261	*target 8 continues*
T	9	0	11.885	8.482	261	*target 9 continues*
T	10	0	1.958	6.538	261	*target 10 continues*
C Thu May 22 10:37:21.484 2003						
T	9	0	14.114	8.290	261	*target 9 continues; target 8 has left*
T	10	0	5.268	6.615	261	*target 10 continues*
T	11	1	15.247	3.869	257	*target11 has just entered with an invalid status*

The column names are the same as used in Figure 4 with an Annotation added to describe the interpretation of the data.

Figure 5. Data Record for Several People in the Field of View

ALGORITHMIC APPROACH

After much experimentation (over a two-year period) the approach we have adopted to tracking people is the following and adopts a multi-pass system because the data is collected and processed in four stages. The first stage is to collect the data from the detectors. All detectors are connected to the same computer to ensure synchronization of the timestamps between the detectors. This process is undertaken in real-time and no further processing is undertaken as the data is collected.

The second stage is to construct, for each target, in each detector, a data structure that comprises for each observed point [x, y, timestamp], obtained from the C and T type records described previously. The number of such points in each data structure varies depending upon the number of times the person is observed as they move through the field of view of the detector. The first point in the sequence of points is referred to as the entry point and the last as the exit point.

For the purposes of further explanation, we shall assume a detector layout as shown in **Figure 6,** which represents a walled corridor. People can enter either from the left of Detector 0 (L) or the right of Detector 2 (R). People can only exit from the same edges. The value of the x-pixel increases as people move from L to R. The next stage of the algorithm is to partition the data structures holding the observed points for each target in each detector into two sorted sequences. One sequence contains all the data structures that have an entry point x-value less than 8; that is, people entering from the left hand side. The other sequence contains all the targets that have entry point x-value grater than 8, that is, all the people entering from the right hand side. The sequences are sorted into time sequence based upon the entry point time. In the experiments reported in the next section, all the data could be fitted into the memory of the processing computer and thus the time to process could be measured in tenths of seconds to undertake this and the last stage.

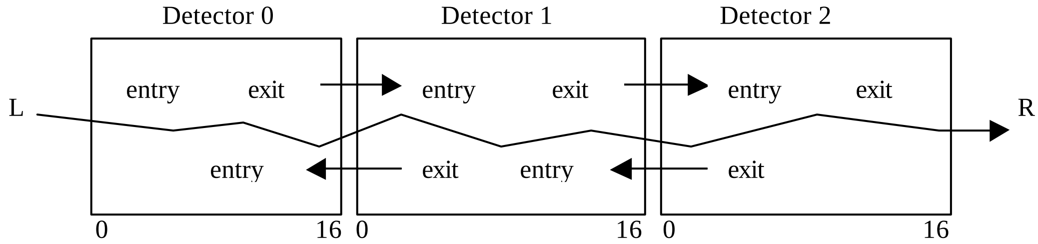

Figure 6. Field of View Relationship

The final stage will be described in terms of the processing required to obtain the trajectories for people entering from the point L and moving towards the point R as shown in **Figure 6**. It will be described assuming that the complete trajectory can be constructed.

```
For each person entering from point L
        Place data structure from Detector 0 into the trajectory
        Obtain corresponding exit point [x, y, timestamp] from data structure
        Fuzzy-match Detector 0 exit point with Detector 1 LHS entry points
        Place data structure from Detector 1 into the trajectory
        Fuzzy-match Detector 1 exit point with Detector 2 LHS entry points
        Obtain corresponding Detector 2 exit point [x, y, timestamp]
        Place data structure from Detector 2 into the trajectory
        The complete trajectory can now be saved to file
```

The above process can then be repeated for people entering from the point R. The `fuzzy-match` operation has to take account of the fact that:

1. detector fields of view may under or over-lap,
2. due to edge effects in the detector a valid data point being matched may be some distance from the edge of the field of view,
3. detector fields of view may not be aligned.

Through the process of experimentation we have determined that using the following value ranges, for each of the above factors respectively can accommodate the fuzzy aspects of the match:

1. ABS (Exit.timestamp – Entry.timestamp) <= 1500 milliseconds
2. Exit.x > 12.5 AND Entry.x < 3.5
3. ABS (Exit.y – Entry.y) <= 2 pixels

Where ABS returns the absolute value of the argument.

These can be interpreted in the following manner

1. If the detectors under or over-lap a person may appear either after or before respectively they have left the detector they are exiting when compared to the one they are entering. The limit of 1500 milliseconds between exiting from one and appearing in the adjacent detector gives a boundary condition that covers both cases.
2. These values ensure that the person is sufficiently close to the edge to be confident that we are observing the same person.
3. The difference in the y-values has to be within 2 pixels, which is equivalent to about 40cm on the ground. Two people when walking together or passed each other, cannot get the center of their torso, viewed from above, within this distance.

If a fuzzy match operation fails for any reason then it will not be possible to determine a complete trajectory for a particular individual.

THE EXPERIMENTAL FRAMEWORK

The experimental framework comprises a corridor 2.5m wide at which we could place the detectors at a height of 3.25m. Three detectors were located down the length of the corridor with the edges of adjacent detectors located just touching in the configuration shown in **Figure 6**.

A group of 28 people were recruited, split into two groups of 14 (referred to as groups A and B) and then asked to undertake a series of movements. These movements were described in global terms such as "the group is to move from L to R". We did not specify how they were to move and whether they were to organize themselves into smaller groups. The movements are summarized in **Table 1** an asterisk indicates an obstruction was placed in the center of the field of view of detector 1.

Move	L		R
1	A → B	→	→
2	← B →	← →	A →
3*	A →	→	→ B
4*	← ←	← ←	A B

Table 1. Sequence of Movements for the Experiment

In **Table 1**, Move 2 corresponds to the 14 members of group A moving from Position R to Position L while at the same time the 14 members of Group B are moving from Position L to Position R.

RESULTS OF THE EXPERIMENT

Table 2 shows the number of people the tracking software was able to identify as they moved from one side of the space to the other according to the specification in **Table 1** in the environment shown in **Figure 6.**

Move	Possible Matches	Direction	Left 0	Boundary 0-1	Boundary 1-2	Right 2	Time secs	% Total Matches
1	14	L to R	13	13	13	13	41	93
2	14	L to R	13	13	13	13	19	93
	14	R to L	11	11	13	13	22	79
3	14	L to R	14	14	13	13	18	93
4	28	R to L	22	22	23	27	20	79

Table 2. Summary Results for the Experiment

The Move column in **Table 2** refers to the Move defined in **Table 1**. The Possible Matches column indicates the total number of matches that are possible for this movement in the direction specified in the Direction column. The Left 0 column gives either the number of people detected entering the field of view of detector 0 if the movement is L to R or the number detected exiting the field of view of detector 0 if the direction is R to L. The Right 2 column gives the same data as the Left 0 column but for the right-hand edge of detector 2. The columns Boundary 0-1 and Boundary 1-2 gives the number of matches that were achieved in the given direction as people moved between adjacent detector fields of view. The final column gives the percentage of people exiting from an edge against the possible number of matches.

Figure 7 shows the trajectories in Move 4 of the 28 people as they moved from Position R to Position L moving around an obstruction placed in the center of the field of view of detector 1. It can be seen how the people deviate round the obstruction from a very close distance and then after the obstruction they spread out to fill the available space.

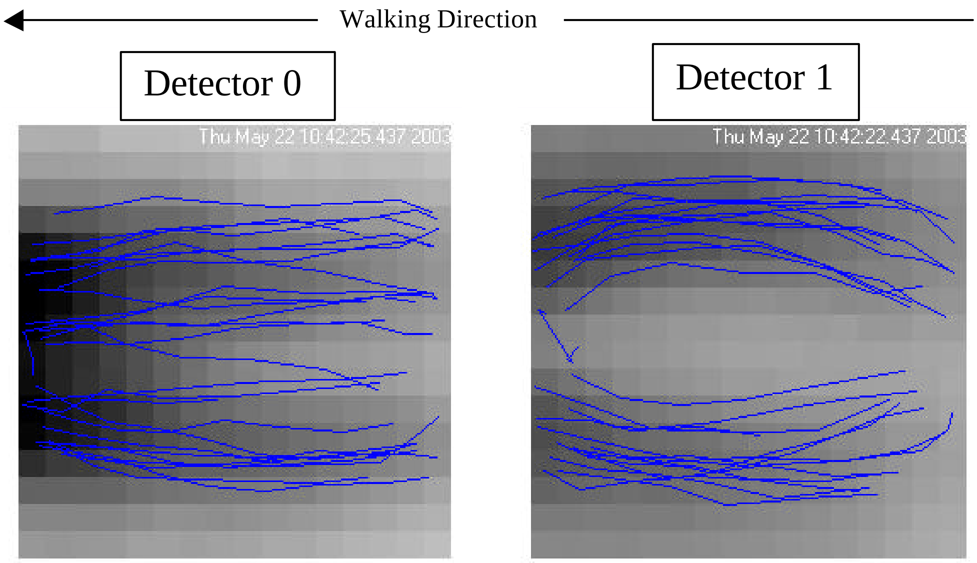

Figure 7. Trajectories of People Moving By an Obstruction

Figure 8 shows the situation in the corridor, where the experiments took place, during Move 3, when 14 people in Group B moved around an obstruction to join the other group at the other side of the experimental area. The obstruction takes the form of a "Caution Slippery" sign, which can be seen in the middle of the picture. The detectors are mounted vertically above the centre of the corridor.

Figure 8. People Undertaking Move 3 Avoiding an Obstruction

EVALUATION

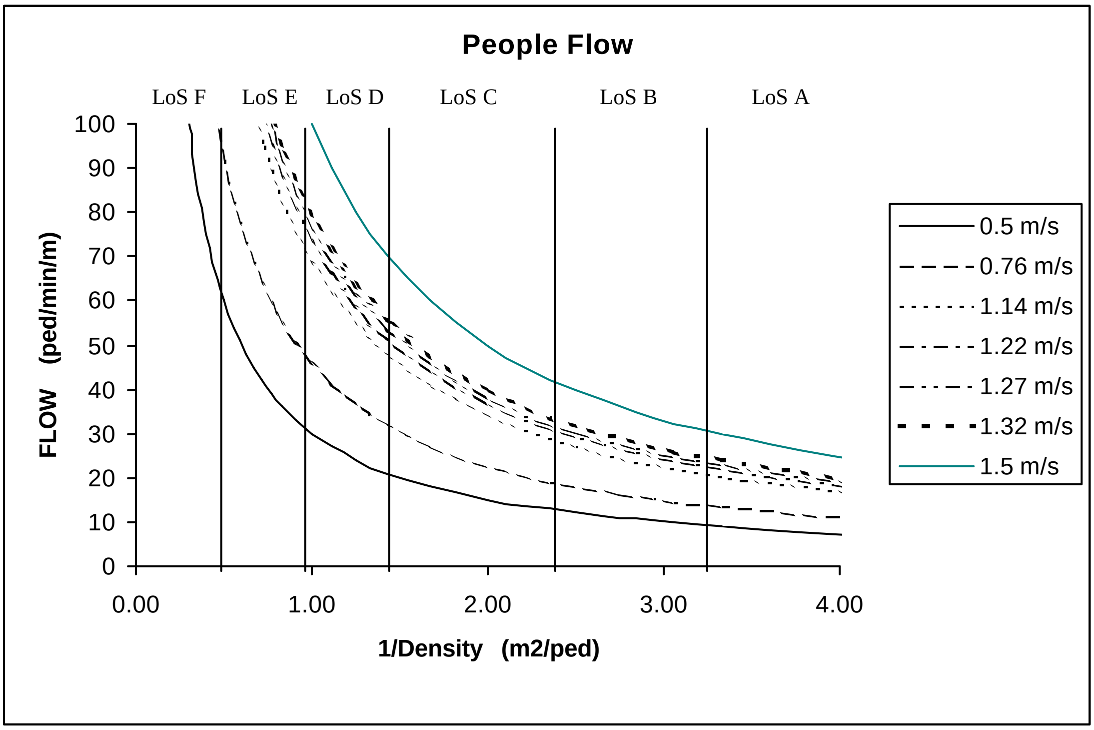

Figure 9. Pedestrian Flow Data

An analysis of the movements in the Experiment based on the work of Fruin [3] using metric values from [24] gives the Density/Flow graph shown in **Figure 9**. The width of the corridor, used for the experiments is 2.26m, which is reduced to 1.34m having removed 0.46m per side to take account of corridor edge effects, as suggested by Fruin. The corresponding flow rates for each movement (people/meter/minute) can then be calculated. The length of corridor over which people are detected is 8.75m. Analysis of the start and end times of the matched individuals for each of the movements yields the average speeds (m/sec) given for each movement. Table 3 summarizes the data and shows the Fruin Level of Service derived from **Figure 9**.

Move	Pedestrian Movements	Time secs	Flow Rate People/m/min	Average Speed (m/sec)	Fruin LoS	% Total Matches
1	14	41	15.29	1.46	A	93
2	14	19	32.99	1.23	C	93
	14	22	28.49	1.33	C	79
3	14	18	34.83	1.36	B	93
4	28	20	62.69	1.31	D	79

Table 3. Evaluation Data for the Experiment

It can be seen that the detectors are able to cope with Fruin service levels A and B reasonably well in non-conflicted flow (Movements 1 and 3 where 93% of the people entering the scene are followed across all the detectors. In Movement 2 where there is conflicted flow at LoS C we can observe some degradation in performance. Once we get to service level D there appears to be further degradation in the system's ability to track an individual over a number of detectors as the number of pedestrians fully tracked drops to 79%. Obviously the system would start to exhibit less good performance if the number of detectors was much larger. The loss of ability to match across a detector boundary stems from the edge effects that are seen in the data and other more radical moves made by individual pedestrians, for example, doing a 'U' turn on the boundary between two detectors. These could be reduced were we to apply more sophisticated algorithms to match trajectories relative to the boundary of adjacent fields of view.

CONCLUSIONS

At the outset of this paper we posed two questions:

- Can we easily track an individual from the field of view of one detector to the next?
- Can we dynamically create datum points over which we can collect various statistics?

Currently, the response to the first would have to be a qualified because further work, as outlined in the next section, has to be undertaken to improve the level of matching that could be achieved. We can be more positive about the second question because even though we could not match every pedestrian across every boundary we were able to count every pedestrian as they entered or exited the area being monitored and return the correct result.

FURTHER WORK

A more detailed set of experiments needs to be undertaken where we can control aspects such as average speeds and flow rates. This will allow us to determine the conditions, in terms of Fruin's Level of Service, where we are able to make effective use of this tracking technology. In particular, we would wish to investigate situations where Levels of Service D and E were apparent. We also are seeking to undertake experiments with various widths of corridors and with people having different levels of mobility so that we can determine the effects these aspects have on the overall throughput of a particular pedestrian space.

We are also investigating techniques whereby we can extract the data from the detectors in real time, rather than using the post-processing techniques we have used in the work reported in this paper. We believe that we will be better able to match pedestrians across detector field's of view because both the location and time interval will be present at the same time and the number of possible matches should be reduced.

Experiments also need to be undertaken to investigate the use of non-linear arrays of detectors. The interactions between the detector processing software becomes much more complex because we have to deal with people moving through the detector fields in an irregular manner.

Finally, we need to undertake experiments in a real urban environment to determine whether the results obtained in the laboratory situation are replicated to areas where we have no control over the behavior of the pedestrians being tracked. We have recently installed the three detectors used for the work reported in this paper at the entrance to one of the buildings of our University. They were installed at a height of 3.5 meters and 3.5 meters apart approximately. The data collection system worked first time and the processing software worked without alteration over many hours of collected data. We shall report on this study in a future paper.

ACKNOWLEDGEMENTS

This research has been supported by a grant from the UK Department of Transport in the LINK programme Future Integrated Transport with the project PERMEATE (GR/N33706). The very helpful comments of the referees are also acknowledged.

REFERENCES

1. Armitage,A., Binnie, T.D., Kerridge, J.M., & Lei,L., (2003), "Measuring Pedestrian Trajectories with Low Cost Infrared Detectors: Preliminary Results", Pedestrian Evacuation and Dynamics – 2003, Galea, E.R. (ed), University of Greenwich, London, UK.
2. Crompton, D.H. and Gilbert, D. (1970), "Traffic and environment - survey and analysis of a group of representative streets", Traffic Engineering and Control, October, pp323-326.
3. Fruin, J.J. (1971), Pedestrian Planning and Design, Metropolitan Association of Urban Designers and Environmental Planners, New York.
4. Goldschmitt, J. (1977), Pedestrian Delay and Traffic Management, Supplementary Report 356, Transport and Road Research Laboratory, Crowthorne.
5. Grayson, G.B. (1975), Observations of Pedestrian Behaviour at Four Sites, Laboratory Report 670, Transport and Road Research Laboratory, Crowthorne.
6. Griffiths, J.D., Hunt, J.G. and Marlow, M. (1984a), "Delay at pedestrian crossings 1. Site observations and interpretation of data", Traffic Engineering and Control, July/August, pp365-371.

7. Griffiths, J.D., Hunt, J.G. and Marlow, M. (1984b), "Delay at pedestrian crossings 2. The development and validation of a simulation model of a zebra crossing", Traffic Engineering and Control, October, pp505-509.
8. Griffiths, J.D., Hunt, J.G. and Marlow, M. (1985), "Delay at pedestrian crossings 4. Mathematical models", Traffic Engineering and Control, May, pp277-282.
9. Hine, J.P. and Russell, J.R.E (1993), "Traffic barriers and pedestrian crossing behaviour", Journal of Transport Geography, Volume 1 (4), pp230-239.
10. Hine, J.P. (1996), "Pedestrian travel experiences: assessing the impact of traffic on behaviour and perceptions of safety using in-depth interview technique", Journal of Transport Geography, Volume 4 (3), pp179-190.
11. Hollock, S,. Galloway, J. L., " Smart Sensors Using Array Based Infrared Detectors", Sensors and their applications XI/ISMCR 2001 - Institute of Physics, City University London, September
12. Hoogendoorn, S.P., Daamen, W. and Bovy, P.H.L., "Extracting Microscopic Pedestrian Characteristics from Video data", Transportation Research Board 2003 Annual Meeting, CD-ROM, Paper No 477, National Academies, Washington USA, 2003
13. Hunt, J.G. and Griffiths, J.D. (1991), Pedestrian Crossing Criteria Research - Random Crossing Model, Contractor Report 248, Transport and Road Research Laboratory, Crowthorne.
14. Hunt, J.G. and Williams, S. (1982), "Delays to pedestrians crossing the road at a random point", Traffic Engineering and Control, April, pp216-221.
15. Mansi, M. V., Porter, S. G. Galloway, J. L. &. Sumpter N., "Very low cost infrared array based detection and imaging systems" (SPIE) Aerosense 2001, Orlando, Florida USA, 17-19 April 2001
16. Navin, F.P.D and Wheeler, R.J.(1969), "Pedestrian flow characteristics", Traffic Engineering, Volume 39, Number 4, pp30-33.
17. Oeding, D. (1963), Traffic loads and dimensions of walkways and other pedestrian circulation facilities, in Strassenbau und Strassenverkehrstechnik, Number 22, Bonn. pp160-163.
18. Older, S.J. (1968), "Movement of pedestrians on footways in shopping streets", Traffic Engineering and Control, August,
19. Reading IAD, Pedestrian Detection Using Computer Vision, PhD Thesis, Napier University, Edinburgh, 1999.
20. Rosin PL and Ellis T, (1995) Image difference threshold strategies and shadow detection, British Machine Vision Conference, 347-356
21. Stickland T and Franks NR, (1994) Computer Image Analysis provides new observations of ant behaviour patterns, Proc R Soc London, Series B, 257,279-286
22. Stogdale, N., Hollock, S., Johnson, N. & Sumpter, N. (2003), "Array based infrared detection: an enabling technology for people counting, sensing, tracking and intelligent detection", SPIE, USE 3, 5071-94
23. Willis, A., Kukla, R., Kerridge, J. & Hine, J. (2001), "Laying the foundations: The use of video footage to explore pedestrian dynamics in PEDFLOW", Proceddings of Pedestrian and Evacuation Dynamics, Shreckenberg,M. et al (eds), Springer, Berlin, pp 181-186.
24. www.crowddynamics.com/Main/Levelsofservice.html (accessed 9/7//2003)
25. www.irisys.co.uk (accessed 9/7/2003)