Improving the Server Performance of CAR Systems Based on Mobile Phones

Víctor Fernández, Juan Manuel Orduña and Pedro Morillo [1]

Resumen— **Collaborative Augmented Reality (CAR) systems allow multiple users to share a real world environment including computer-generated images in real time. The hardware features of most current mobile phones include wireless network capabilities that offer a natural platform for CAR systems. However, the potential number of clients in CAR systems based on mobile phones is much larger than on CAR systems based on other kind of mobile devices, requiring a system design that takes into account scalability issues. This paper presents the experimental comparison of different CAR systems based on mobile phones with different server implementations. The performance evaluation results show that the best implementation is the one based on UDP messages instead of classical TCP connections, in order to improve the system throughput. The UDP-based implementation provides a significant improvement in system throughput, at the cost of loosing a very small percentage of updating messages. However, the effects of these small quantities of dropped messages cannot expand beyond some jitter (bounded within a short period of time) in a reduced number of clients of the CAR application. These results validate the proposed UDP-based implementation as the best option for large-scale CAR systems based on mobile phones.**

Palabras clave— **Collaborative Augmented Reality, Mobile Phones.**

I. Introducción

AUgmented Reality (AR) systems are nowadays widely used in applications such as medical procedures, scientific visualization, manufacturing automation, cultural heritage and military applications. The term Augmented Reality (AR) refers to computer graphic procedures or applications where the real-world view is superimposed by computer-generated objects in real-time [1], [2]. From the beginning of AR systems, the potential of collaborative AR (CAR) systems was exploited for different activities such as Collaborative Computing or Teleconferencing [3]. Wearable devices were used to provide CAR systems, where a wearable AR user could collaborate with a remote user at a desktop computer [4]. On other hand, mobile phones have become an ideal platform for CAR systems, due to the multimedia hardware that they include. As an example, Figure 1 shows a CAR system developed for collaborative training in industrial electricity. It shows on the left image the execution of the CAR tool on a Samsung Galaxy NOTE mobile phone, while the image on the center shows a real image of the the panelboard where technicians collaboratively operate, and the right image shows the execution of the CAR tool on a HTC Nexus One mobile phone.

Fig. 1. Example of a CAR application developed for training in industrial electricity.

The wide variety of current mobile phones, with different graphic and processing capabilites, and different operating systems, can have significant effects on the performance of a large-scale CAR system, in terms of system latency, frames per second or number of supported clients with certain latency levels. In previous works, we have characterized the behavior of different mobile phones and the server when used in Collaborative Augmented Reality applications, [5], [6]. The results showed that CAR systems throughput heavily depends on the kind of client devices, but for certain kind of devices, the system bottleneck is the server I/O.

In this paper, we propose a comparative study of different implementations of the CAR server, in order to improve the performance of CAR systems based on mobile phones. The performance evaluation results show the UDP-based implementation provides a significant improvement in system throughput with respect to other implementations, supporting more than one thousand clients at interactive rates (twice the number of supported clients of the TCP implementation). This improvement is achieved at the cost of loosing a very small percentage of updating messages but the effects of these dropped messages cannot expand beyond some jitter (bounded within a short period of time) in a reduced number of clients.

The rest of the paper is organized as follows: Section II shows some related work about CAR applications on mobile phones. Section III describes the different CAR implementations considered for comparison purposes, and Section IV shows the performance evaluation results. Finally, Section V presents some conclusion remarks.

The rest of the paper is organized as follows: Section II shows some related work about CAR applications on mobile phones. Section III describes the different CAR implementations considered for comparison purposes, and Section IV shows the perfor-

---

[1] Dpto. de Informática, Univ. of Valencia, Valencia, Spain. E-mails: {`Victor.Fernandez-Bauset, Juan.Orduna,Pedro.Morillo`}`@uv.es`

mance evaluation results. Finally, Section V presents some conclusion remarks.

II. Related Work

Augmented Reality superimposes multimedia content - 3D object, text, sound, etc - to real world through a display or screen. In order to locate digital contents on a specific image of the real world point, some references within the image are needed. These references are known as markers, and two methods are usually used: natural feature tracking and fiducial marker tracking. The former method uses interest point detectors and matching schemes to associate 2D locations on the video with 3D locations [7]. This process can be grouped in three big phases: interest point detection, creation of descriptor vectors for these interest points, and comparison of vectors with the database [8]. The latter method uses fiducial markers to find a specific position of real world. This process can be divided in three phases: edge detection, rejection of quadrangles that are too large or too small, and checking against the set of known patterns [7].

Any CAR application needs a device equipped with an on-board camera, CPU and display. The most common devices used for CAR applications are Tablet-PCs or mobile phones. We will focus on mobile phones, because they are more suitable for CAR applications [9].

There are few solutions based on fiducial marker tracking over mobile phones. In 2003, ArToolKit [10], one of the most well-known software libraries for developing Augmented Reality (AR) application, was released for Windows CE, and the first self-contained application was developed for mobile phones [11]. This software evolved later as the ArToolKitPlus tracking library [7]. A tracking solution for mobile phones that works with 3D color-coded marks was developed [12], and a version of ArToolKit for Symbian OS was developed, partially based on the ArToolKitPlus source code [13]. The research teams behind these works have worked on fiducial marker tracking, but not from the collaborative point of view. Also, there are many other works that focus on natural feature tracking [7], [14], [15], [16].

Although real-time natural feature tracking over mobile devices has been currently achieved [7], fiducial marker tracking is more widely used, because it allows simultaneous computational robustness and efficiency. A large number of locations and objects can be efficiently labeled by encoding unique identifiers on the markers. Additionally, the markers can be detected with angles near to 90 degrees [7].

The first CAR applications improved the conference system highlights, giving the feeling of real presence to remote collaborators [3]. The Rekimoto's Transvision system showed how to share virtual objects through handheld displays [17]. Also, Schmalstieg created a software architecture to develop CAR applications [18].

III. Server implementations

We have developed a multithreaded CAR server that supports simulated clients (simulated mobile devices) with the behavior measured in our previous work [5]. The system configuration consists of this server and a certain amount of mobile devices that are scanning the visual space of their video camera looking for a marker that will be converted into a 3D object in their display. After each updating of the object location, the mobile device sends a location update message (containing the new location) to each of its neighbor devices. The neighbor devices are those who participate in the same collaborative task, and we have denoted this set of neighbor devices as a *working group*. The messages are sent through the server (that is, it sends the location update message to the server, and then the server re-sends the message to the appropriate clients). For performance evaluation purposes, the destination clients return an acknowledgment message (ACK) to the server, which, in turn, forwards it to the source client. When the source client has received the ACK messages corresponding to the location update from all the clients in its working group, then it computes the average system response for that location update. Figure 2 illustrates the action cycle that takes place for each of the mobile clients in the system.

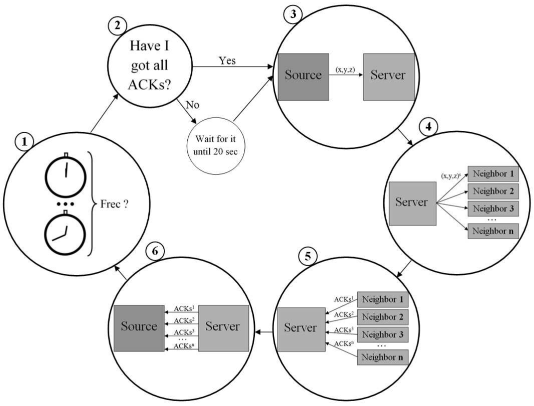

Fig. 2. Stages of the action cycle in each mobile device.

Once the message with the location update is sent, the action cycle performed by each client is composed of the following steps: first, it performs one new image acquisition followed by a marker detection stage. Then, the client waits until the cycle period (determined by the action frequency, a system parameter) finishes. Next, if the acknowledgments from all the neighbors have been received, a new message with the new marker location is sent. If not all the acknowledgments have been received, then it waits until a maximum threshold of 20 seconds, and then a new round of messages (with the latest marker location) are sent to the neighbors through the server. The neighbors simply return an ACK message to the sender device through the server. The server simply forwards the messages to the corresponding destination clients. It must be noticed that the mobile devices will not send a new round of messages with a new location update until it has received the acknowledgment message from all its neighbors, even although new marker detection stages have been completed in the device.

This characterization setup considers that all the required static content in the scene has been loaded. According to recent works [19], in these cases the network bandwidth required is less than 50 kbps for performing this information exchange. Since we are using a Gigabit Ethernet, we ensure that the network bandwidth does not become a system bottleneck.

The system latency provided for each location update is computed by recording a timestamp when the first message is sent to the server. Next, a second timestamp is recorded with the last ACK message for that location update received from the server. The system response time is computed by subtracting these two timestamps. The server response time is computed by timestamping both each message forwarded from each client and the reception of the corresponding ACK message from the destination client. Also, the percentage of CPU utilization is measured both in the server and the mobile devices every half second.

A. TCP Implementation

The simulator starts generating a *Server Process*, and for every 50 clients it generates a *Client Process*. Figure 3 illustrates the general scheme of the Server Process. This process starts listening connections, and for each connection it generates a new array of X TCP sockets, where X is the number of clients that will be within a given working group. When all the clients have connected to the Server Process (the population size is a simulation parameter) then the Server Process generates as many *Server Threads* as needed. Each Server Thread is in charge of managing all the clients within a working group. Concretely, it starts the simulation by sending a welcome message to all the client sockets. When the simulation finishes, it collects statistics from all the clients in its working group. But the most important task performed by server threads is the generation of two threads for each of the clients in the working group: the *Server Receiver Thread (SRT)* and the *Server Processor Thread (SPT)*. The SRT associated to client i receives the location update messages from the client i. Next, it computes the correct destination clients (the neighbor clients, that is, the clients within the same working group) and it generates messages that will be stored in the queues of the Server threads managing these neighbor clients. The SPT associated to client i extracts the queued messages that the SRTs associated to other clients may have generated for client i, and it sends them to this client. Additionally, the server process collects and processes the statistics generated by the server threads, and it also measures the percentage of CPU utilization.

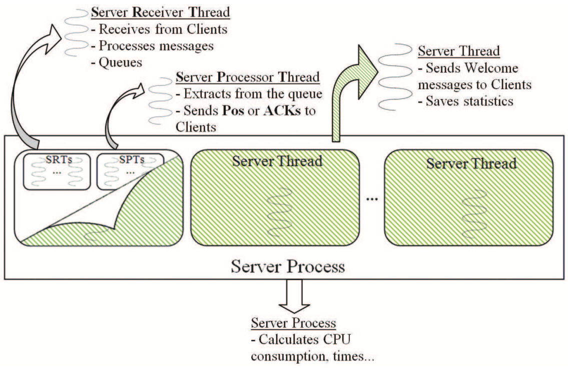

Fig. 3. General scheme of the server process in the TCP implementation.

Figure 4 illustrates the general scheme of the Client Process. This process generates 50 client threads (we have assumed a maximum population size of 1000 client devices), and it also computes the percentage of CPU utilization, client latencies, etc.. Each Client Thread generates two threads for each client: the *Client Receiver Thread (CRT)* and the *Client Processor Thread (CPT)*, and when the welcome message from the Server Thread arrives to the associated socket, then the Client Thread starts the simulation, that consists of sending a given number of position update messages and receiving the corresponding acknowledgments from the neighbor clients.

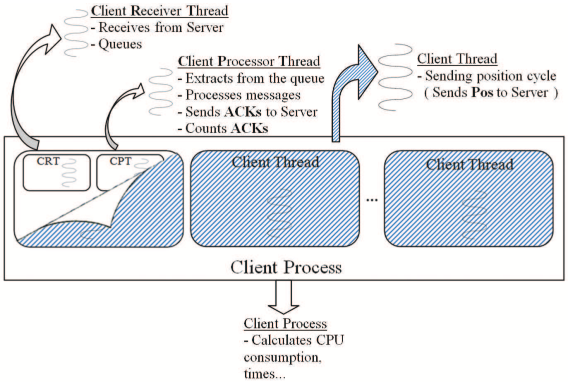

Fig. 4. General scheme of the client process in the TCP implementation.

Also, we developed another version where each Server Thread has a single SRT and a single SPT for managing all the clients in each working group, instead of one SRT and one SPT for each client. Using the `Select` function, the SRT receives messages from all the clients and it processes them. As it could be expected, we obtained better performance results with the Select version of the TCP implementation.

B. UDP Implementation

Finally, we have considered a connectionless oriented implementation for the CAR system, in order to study the effectiveness of TCP connections in a distributed environment like a CAR system. The motivation of this study are both the short message size (usually carry a position update consisting of a bunch of bytes) and the huge amount of the messages generated by CAR systems. Although the UDP protocol can loose messages and the effects and size of these losses should be studied, we have also considered this implementation for comparison purposes. The UDP implementation is very similar to the TCP-Select implementation. The only difference is that in this implementation we have used UDP sockets. Since this implementation can drop messages, it also counts the number of dropped or lost messages (since both the number of iterations and the number of clients in each working group is known, each client can compute the number of message that should arrive).

IV. Performance Evaluation

We have performed different measurements on different simulated systems using these implementations. We have performed simulations with different number of clients and we have measured the response time provided to these clients (the round-trip delay for each updating message sent by a given client to the clients in its working group). In this way, we can study the maximum number of clients that the system can support while providing a response time below a given threshold value. In order to define an acceptable behavior for the system, we have considered 250 ms. as the threshold value, since it is considered as the limit for providing realistic effects to users in DVEs [20].

We have considered the system response time (in milliseconds) for each updating message sent by a given client to its neighbor clients as the time required for receiving the acknowledgments from all the clients in the working group of this given client. In order to measure the dispersion of this metric, we have measured the standard deviation for all the updating messages sent, as well. Also, we have computed the response time in the server (in milliseconds) as the time required by the destination clients to answer the server messages. We have measured both the average and the maximum values measured in the server for each simulation. Additionally, we have computed the percentage of the CPU utilization in the system server, since it can easily become the system bottleneck. The computer platform hosting the system server is a Intel Core 2 Duo E8400 CPU running at 3.00 GHz with 4 Gbytes of RAM, executing an Ubuntu Linux distribution with the 3.0.0-14-generic x86 64 operating system kernel. In order to study the system behavior for different levels of workload, we have repeated simulations with working group sizes of 10,15,20 and 25 clients. Due to space limitations, we only show here the results for the biggest size (25 clients in each working group).

Table I shows the results for a CAR system whose client devices are all of them HTC Nexus One, and where the working group size for each client is of five neighbor clients. This table shows the results for the three considered implementations, organized as three subtables with ten rows each, and labeled with the name of the implementation (TCP, TCP-Select and UDP). The most-left column in these subtables shows the number of clients in the system, that is, the population size. The values in this column range from 100 to 1000 clients in the system. The next two columns show the average value of the response times (in milliseconds) provided by the system to all the clients (labeled as "RT"), as well as the corresponding standard deviation values (column labeled as "Dev"). The fourth column (labeled as "CPU") shows the percentage of the CPU utilization in the server. Finally, the fifth and sixth columns (labeled as "RT_S" and "RT_SM", respectively) show the average and maximum values (in milliseconds) of the response time in the server for all the messages exchanged during the simulation.

	TCP implementation				
Size	RT	Dev	CPU	RT_S	RT_SM
100	62.37	22.68	9.9	19.36	20.9
200	63.77	22.21	15	20.12	22.43
300	66.71	22.66	22	25.15	28.28
400	68.68	22.5	32.7	27.14	29.56
500	71.04	23.56	45	27.14	30.9
600	71.5	24.18	48.6	26.58	31.95
700	72.37	25.01	59	27.17	35.42
800	72.85	26.01	68	27.87	34.54
900	75.01	28.98	79.2	28.61	35.89
1000	147.33	101.71	85	43.95	48.66

	TCP-Select implementation				
Size	RT	Dev	CPU	RT_S	RT_SM
100	65.77	21.56	8	20.44	24.96
200	67.12	22.71	11.2	21.5	23.54
300	67.52	22.6	19.8	23.67	26.62
400	67.64	22.88	28	23.21	27.28
500	69.12	23.23	31	25.31	29.34
600	69.14	23	39.6	25.28	28.74
700	69.37	23.45	47	24.53	30.72
800	75.75	26.63	54.5	26.96	35.16
900	70.24	24.81	59.6	24.02	31.87
1000	71.05	27.53	67	22.64	35.04

	UDP implementation				
Size	RT	Dev	CPU	RT_S	RT_SM
100	4.80	7.06	38.40	1.95	3.22
200	3.76	4.95	34.70	1.57	2.84
300	9.57	9.74	26.00	4.44	10.97
400	3.60	5.18	33.70	1.46	3.16
500	4.59	6.41	41.60	1.81	3.77
600	7.34	13.17	47.00	3.23	17.16
700	5.28	8.24	53.00	2.11	7.76
800	7.10	18.52	84.10	2.65	16.53
900	5.85	11.69	66.30	2.61	12.05
1000	7.15	15.47	69.50	2.87	15.18

TABLA I

Results for a working group size of 5 neighbors

Table I shows that none of the values in the RT column reaches the threshold value of 250 milliseconds in any of the considered implementations, showing that the system can efficiently support up to one thousand clients while interactively displaying the Augmented Reality content. Nevertheless, there are significant differences in this column among the considered implementations. Thus, the TCP implementation shows a huge rise in the response time when the system reaches one thousand clients, passing from around 75 milliseconds to more than 147 milliseconds as an average. The standard deviation of these values are also more than three times the value shown for nine hundred clients. These values show that for that population size the system is approaching saturation. On the contrary, the TCP-Select implementation does not show an increase in neither the column RT nor the column Dev for a population of one thousand clients. Moreover, the UDP implementation shows RT values that are one order of magnitude lower than the ones shown by the other two implementations.

The third column in table I shows that the CPU utilization increases as the number of clients in the system increases. In the case of the TCP implementation, the system approaches saturation when the server reaches 85% of CPU utilization. For lower percentages of CPU utilization the response times do not significantly increase. It is worth mention that the UDP implementation provides RT values that are one order of magnitude lower than the ones provided by the TCP implementations, even for CPU utilization of around 70%. These values show that the latency provided by CAR systems greatly depends on the connection or connectionless scheme followed by the system to exchange information with the clients.

Finally, the columns RT_S and RT_SM show that most of the response time provided to clients is due to processing in the server. Thus, for example, the results for the TCP implementation and a system size of 900 clients show that, as an average, each client has to wait 75.01 milliseconds for receiving the acknowledgments from all the clients in its working group, but as an average the server must wait only 28.61 milliseconds to receive answers from clients. This difference highly increases for the case of one thousand clients, where the response time obtained by the server from clients is around 44 milliseconds but the average response time provided to clients is 147.33 milliseconds, around three times higher. It is also worth mention that the ratio between the RT_S and the RT columns do not significantly vary among the three implementations. Finally, the RT_SM column shows that the maximum values in the RT_S parameter do not exceed the value in the RT column for both TCP implementation, and they do not exceed twice the value in the RT column of the UDP implementation. Therefore, we can conclude that most of the time required to acknowledge each client update is due to the processing of the updates and acknowledgments in the server.

These results show that the best latencies when the system is far from saturation are provided with the UDP implementation. However, UDP is a connectionless-oriented protocol, and therefore it may drop messages when the system approach saturation.

Table II shows the results for a working group size of 25 clients. The most-left column in these subtables shows the number of clients in the system, that is, the population size. The values in this column range from 100 to 1000 clients in the system. The next two columns show the average value of the response times (in milliseconds) provided by the system to all the clients (labeled as "RT"), as well as the corresponding standard deviation values (column labeled as "Dev"). The fourth column (labeled as "CPU") shows the percentage of the CPU utilization in the server. Finally, the fifth and sixth columns (labeled as "RT_S" and "% lost", respectively) show the average values (in milliseconds) of the response time in the server for all the messages exchanged during the simulation and the percentage of messages dropped by the system. The latter column has been computed by subtracting the number of messages received by all the clients in a simulation (measured in the simulation itself) from the theoretical number of messages that clients should exchange for a given population size.

	TCP-Select implementation				
Size	RT	Dev	CPU	RT_S	% lost
100	90.8	24.7	23.2	19.35	0.00
200	89.95	21.13	47	33.4	0.00
300	123.95	32.36	72	54.7	0.00
400	209.2	35.88	87.2	85.55	0.00
500	268.17	44.44	86	112.07	0.00
700	383.96	70.6	93.1	151.56	0.00
1000	566.44	133.33	93.1	166.79	0.00

	UDP implementation				
Size	RT	Dev	CPU	RT_S	% lost
100	9.86	6.78	72.50	4.06	0.83
200	21.70	14.73	82.00	9.84	1.18
300	26.01	21.91	79.60	11.61	0.69
400	39.41	30.66	81.90	18.26	0.83
500	48.68	39.68	83.80	22.84	0.74
700	79.70	97.87	85.10	37.26	0.76
1000	122.37	85.35	85.00	44.98	0.90

TABLA II

RESULTS FOR A WORKING GROUP SIZE OF 25 NEIGHBORS

Table II shows that for this level of workload the system enters saturation in the TCP-based implementation. Effectively, the RT column shows that TCP-Select implementation reaches (and exceeds) this threshold value for a population of 500 clients. However, the UDP implementation does not reach even half of this value for the maximum population size considered, one thousand clients.

It is worth mention that for those cases when the system reaches saturation, the percentage of CPU utilization in the server is 85% or higher. The gap between 85% and 98% of CPU utilization for reaching the saturation point can be explained by the shared memory architecture of current multicore processors (the dual core processor in the computer platform used as simulation server), as shown in [6].

The "% loss" column shows that for the UDP implementation the percentage of lost messages is not higher than 1.2%. The effects of loosing some messages will consist of producing some jitter in the display of the clients. However, these percentage values ensure a reasonable quality in the visualization of the CAR system. In order to ensure that the effects of the UDP implementation in terms of dropped messages are consistent for all the workload levels considered, Figure 5 shows the average number of packets dropped for each working group size considered.

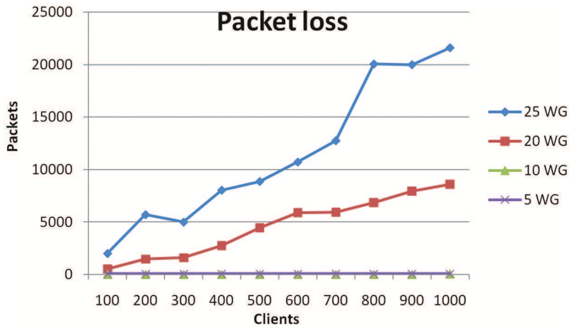

Fig. 5. Number of packets lost in the UDP implementation.

Figure 5 shows that for working group sizes of 5 and 10 neighbor clients there are no packet losses. For a working group size of 20 neighbors, the amount of lost packets reaches 8581 for a theoretical total of packets sent of 1.9 million packets. Analogously, for a working group size of 25 neighbors, the amount of lost packets reaches 21593 out of 2.4 million packets sent. Therefore, in the worst case the number of lost packets only represent a 1'18 % of the total amount of packets sent. This value represents only a small image flicker on some clients, and in very limited periods of time. As the information is sent more than once per second (since the action cycle of the HTC Nexus One is 167.11 ms.), this value can be considered an insignificant flickering.

Although they are not here for the shake of shortness, we repeated the same tests shown in this section here using a different client device, the Motorola Milestone, and we obtained analogous results. Those results were less interesting because of the bigger action cycle of the Milestone (698.34 ms.). With that action frequency the system saturation point was not reached even in the worst case of a working group size of 25 neighbors and a population of 1000 clients. We have shown here the results for the Nexus One as the worst case for the server implementation.

V. CONCLUSIONS

This paper has proposed the experimental comparison of different large-scale CAR systems based on mobile phones with different server implementations. The performance evaluation results show that the best implementation is the one based on UDP messages, instead of classical TCP connections, in order to support more than one thousand clients at interactive rates. These results validate the proposed UDP-based implementation as the best option for large-scale CAR systems based on mobile phones.

ACKNOWLEDGEMENTS

This work has been jointly supported by the Spanish MICINN and the European Commission FEDER funds, under grant TIN2009-14475-C04.

Referencias

[1] R. Azuma, Y. Baillot, R. Behringer, S. Feiner, S. Julier, and B. MacIntyre, “Recent advances in augmented reality,” *Computer Graphics and Applications, IEEE*, vol. 21, no. 6, pp. 34 –47, 2001.

[2] S. Cawood and M. Fiala, *Augmented Reality: A Practical Guide*, Pragmatic Bookshelf, 2008.

[3] M. Billinghurst and H. Kato, “Real world teleconferencing,” in *Proc. of the conference on Human Factors in Computing Systems (CHI 99)*, 1999.

[4] Tobias Hallerer, Steven Feiner, Tachio Terauchi, Gus Rashid, and Drexel Hallaway, “Exploring mars: Developing indoor and outdoor user interfaces to a mobile augmented reality system,” *Computers and Graphics*, vol. 23, pp. 779–785, 1999.

[5] Victor Fernández Bauset, Juan M. Orduña, and Pedro Morillo, “Performance characterization on mobile phones for collaborative augmented reality (car) applications,” in *Proceedings of the 2011 IEEE/ACM 15th DS-RT*, 2011, DS-RT '11, pp. 52–53.

[6] Victor Fernández Bauset, Juan M. Orduña, and Pedro Morillo, “On the characterization of car systems based on mobile computing,” in *Proceedings of HPCC '12 (AH-PCN workshop)*, 2012.

[7] Daniel Wagner, Gerhard Reitmayr, Alessandro Mulloni, Tom Drummond, and Dieter Schmalstieg, “Pose tracking from natural features on mobile phones,” in *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*, Washington, DC, USA, 2008, ISMAR '08, pp. 125–134, IEEE Computer Society.

[8] Seung Eun Lee, Yong Zhang, Zhen Fang, S. Srinivasan, R. Iyer, and D. Newell, “Accelerating mobile augmented reality on a handheld platform,” in *Computer Design, 2009. ICCD 2009. IEEE International Conference on*, October 2009, pp. 419 –426.

[9] Michael Haller; Mark Billinghurst; Bruce Thomas, *Emerging Technologies of Augmented Reality: Interfaces and Design*, IGI Global, 2007.

[10] Dr. Hirokazu Kato, “Artoolkit,” 2011, Available at http://www.hitl.washington.edu/artoolkit/.

[11] Daniel Wagner and Dieter Schmalstieg, “First steps towards handheld augmented reality,” in *Proceedings of the 7th IEEE International Symposium on Wearable Computers*, Washington, DC, USA, 2003, ISWC '03, pp. 127–135, IEEE Computer Society.

[12] Mathias Mahring, Christian Lessig, and Oliver Bimber, “Video see-through ar on consumer cell-phones.,” in *ISMAR'04*, 2004, pp. 252–253.

[13] A. Henrysson, M. Billinghurst, and M. Ollila, “Face to face collaborative ar on mobile phones,” in *Mixed and Augmented Reality, 2005. Proceedings. Fourth IEEE and ACM International Symposium on*, October 2005, pp. 80 – 89.

[14] S. Srinivasan, Zhen Fang, R. Iyer, S. Zhang, M. Espig, D. Newell, D. Cermak, Yi Wu, I. Kozintsev, and H. Haussecker, “Performance characterization and optimization of mobile augmented reality on handheld platforms,” in *Workload Characterization. IISWC 2009. IEEE International Symposium on*, 2009, pp. 128 –137.

[15] D. Wagner, G. Reitmayr, A. Mulloni, T. Drummond, and D. Schmalstieg, “Real-time detection and tracking for augmented reality on mobile phones,” *Visualization and Computer Graphics, IEEE Transactions on*, vol. 16, no. 3, pp. 355 –368, May-June 2010.

[16] D. Wagner, D. Schmalstieg, and H. Bischof, “Multiple target detection and tracking with guaranteed framerates on mobile phones,” in *Proceedings of ISMAR 2009*, 2009, pp. 57 –64.

[17] Jun Rekimoto, “Transvision: A hand-held augmented reality system for collaborative design,” in *Virtual Systems and Multi-Media (VSMM)'96*, 1996.

[18] Z. Szalavari, D. Schmalstieg, A. Fuhrmann, and M. Gervautz, “študierstube̤: An environment for collaboration in augmented reality,” *Virtual Reality*, vol. 3, pp. 37–48, 1998.

[19] Tuomas Kantonen, “Augmented collaboration in mixed environments,” M.S. thesis, Helsinky University of Technology, 2009.

[20] T. Henderson and S. Bhatti, “Networked games: a qos-sensitive application for qos-insensitive users?,” in *Proceedings of the ACM SIGCOMM 2003*. 2003, pp. 141–147, ACM Press / ACM SIGCOMM.