CMUseum: A Location-aware Wireless Video Streaming System[1]

Mei-Hsuan Lu and Tsuhan Chen

Electrical and Computer Engineering Department
Carnegie Mellon University
{meihsual, tsuhan}@ece.cmu.edu

ABSTRACT

Location-aware computing enables automatic tailoring of information and services based on the current location of a mobile user. We have designed and implemented CMUseum, a system that enables location-based video streaming applications, as well as other add-on services on top of the purposed infrastructure. Our novel design incorporates an 802.15.4 Zigbee sensor network for collecting location information as well as an 802.11 Wi-Fi network for streaming video contents. This framework is largely event-driven in order to support the real-time nature of the video streaming services it handles. We have demonstrated the viability of this system through implementing a working system in a museum tour guide scenario.

1. INTRODUCTION

Over the past few years researchers at Carnegie Mellon University have been developing new sensing and multimedia technologies for location-aware computing. This technology can be used for many location-aware services, including traffic flow control and disaster monitoring, city parking assistance and planning, community nursing, warehouses inventory tracking, surveillance, industrial control, etc. In this paper, a practical multimedia sensing system is presented. We call it *CMUseum* since this system has been successfully deployed in a museum touring scenario. CMUsuem provides location-aware video streaming services with other add-on features which enable automatic tour guidance without user intervention. However, the potential applications are not restricted to the museum touring system and can be extended to a host of different location-computing multimedia services.

Location-aware computing has been studied by many researchers [1][2][3]. In [1], the authors designed and implemented *Rover*, a system that enables location-based services as well as the traditional time-aware, user-aware, and device-aware services. The underlying software architecture has also been validated through the initial implementation. In [2], Wang introduces *Nidaros*, a framework for developing location-aware applications that provide location dependent functionality based on the current location of the user is presented. In [3], Chen *et al.* present an infrastructure that supports location-aware services. This infrastructure is based on a proposed location operating reference model, *LORE*, which addresses many major respects of building location-aware services.

Our work differs from the above efforts in two fundamental respects: First, the location information and video contents are transmitted in two wireless networks operating on different frequency bands. In particular, the real-time location information is transmitted with an 802.15.4 Zigbee [4] sensor network running with a real-time TDMA-based MAC protocol, while the video data are transmitted with a random access 802.11 Wi-Fi [5] network. Compared to other positioning technology like GPS, GSM, Bluetooth, and IR, the Zigbee sensing infrastructure is low cost, easy to deploy, and provides better precision. Despite the low bandwidth nature of 802.14.5 network, its capacity is sufficient for sending location information in real-time, low power consumption fashion. The 802.11 network fulfills the large bandwidth requirements for transmitting video data.

Secondly, in our design of CMUsem, the video sever provides more sophisticated streaming services which utilize location information in mapping the networking parameters. Specifically, the location information is used as an index to the transmission rate

in 802.11 networks. The video server can therefore transcode or select the appropriate video stream to fit the available bandwidth. The video server also considers the number of concurrent users to ensure acceptable video quality for each user.

The rest of this paper is organized as follows: Section 2 describes the system architecture which gives the readers an overview of CMUseum. Section 3 presents the sensor network technology adopted in the framework. Section 4 shows the video streaming system. Section 5 introduces other add-on services developed on top of the proposed infrastructure. Finally, Section 6 suggests future improvements and Section 7 reflects upon our conclusions.
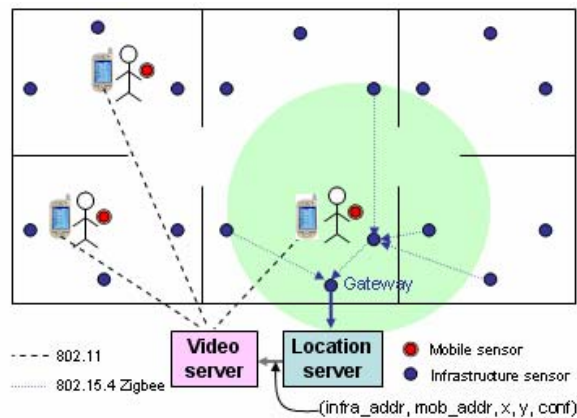
2. SYSTEM ARCHITECTURE

Figure 1. System architecture

The system architecture of CMUseum is shown in Figure 1. This museum has been laid out with a network of wireless infrastructure sensors (Figure 2(a)). In addition, a mobile sensor is attached with each portable display device to enable the location tracking of a particular user (Figure 2(b)). Each sensor has a radio transmitter and is capable of communicating with its neighbors, from hop to hop and eventually to a local gateway. The infrastructure nodes report the visitor's presence to their neighbors which in turn forward the information across one or more hops to the gateway. The gateway then triangulates and localizes the position of the visitor from multiple sensor readings. The sensor network reports every user's location to the gateway and provides timeliness guarantees for end-to-end communication through the 802.15.4 Zigbee physical layer. The gateway then relays mobile node locations to the video server. Using this information, whenever there is a location change for a particular user, the video server streams a different video sequence which corresponds to the user's new location.

By proper placement of the infrastructure nodes and power control for the mobile nodes, we can assure a good precision in locating the end user.
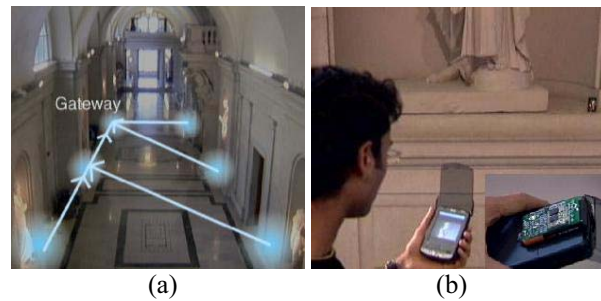

Figure 2. (a) Zigbee sensor network (b) museum touring scenario

The gateway node interfaces with the location server using an external USB dongle. Once a location change event is determined, the location server uses a TCP socket connection to deliver the closest infrastructure node's MAC address, the mobile node's MAC address, the coordinate and the confidence level of the new location to the video server. The video server then finds the corresponding IP address of the PDA and adjusts the video content accordingly. The video streaming mechanism is designed for operation over wireless channels with high packet loss. The 802.11 Wi-Fi [1] technology is chosen as the video streaming network because it is widely available, low price, and easy to configure on hand held devices. This framework is largely event-driven in order to support the real-time nature of the location information it handles.

In the proposed infrastructure, we adopt two different networking technologies: 802.15.4 Zigbee and 802.11 Wi-Fi, for collecting location information and streaming video contents. By decoupling the occupied frequency bands with two physical networks, we can minimize the chance of interference and optimize the performance of each network independently. This low-cost architecture delivers many capabilities while requiring a manageable level of complexity (particularly in terms of implementation and deployment). In the following sections, we describe the sensor network and video streaming system, respectively.

3. SENSOR NETWORK TECHNOLOGY

For the sensor network, we develop a time-synchronized link layer protocol, called *RT-Link* [6], to ensure a collision-free and energy-efficient multi-hop wireless networks. RT-Link facilitates dynamic

admission of both infrastructure and mobile nodes into a tightly synchronized regime. It schedules nodes in time slots such that concurrent transmitters do not interfere with each other and the activity of all nodes are coordinated to maximize the sleep duration. RT-Link maintains contention-free operation by employing an online and automatic link conflict detection and resolution scheme. Such a scheme is preferable in serving real-time or delay sensitive applications. In the current implementation of CMUseum, we assure the location information can be determined in less than one second. RT-Link has been implemented as a link layer protocol in low-cost and low-power embedded nodes developed by us. Interested readers can refer to [6] for details of this technology. Figure 3 shows the hardware boards of the infrastructure and mobile sensor nodes. The daughter board attached on top of the infrastructure node is a component used for global time synchronization.
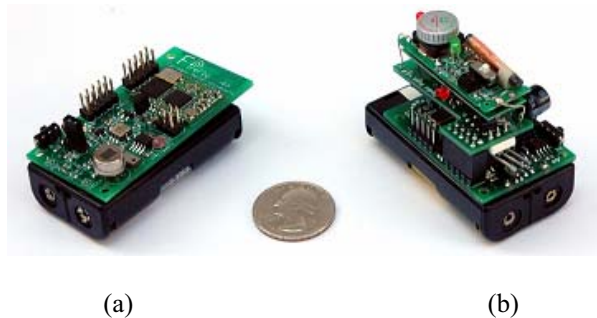

(a)   (b)

Figure 3. (a) mobile node (b) infrastructure node

4. VIDEO STREAMING SYSTEM

The video streaming server is developed based on VideoLAN [7] VLC 0.8.1. A WinCE/PocketPC-based VLC 0.8.3 is running on each PDA, serving as the video client. VLC is cross-platform, highly portable media player for various audiovisual formats as well as various streaming protocols. Given a visitor moving with walking speed and the statues are placed with intervals of several meters, we have modified VLC's control flow to ensure that video content switching is achievable in real time.

The location information enables the video streaming system to provide more sophisticated video services. In the 802.11 specification, rate adaptation schemes that adjust to varying channel conditions have been used in order to enhance the throughput and channel utilization of a wireless link. In general, sender/receiver pairs separated by longer distances suit lower transmission rates, and vice versa. Unfortunately,

due to protocol layering, such MAC layer information is not accessible to the upper layers. To facilitate video processing in the application layer, researchers have proposed different mechanisms such as bandwidth estimation, feedback channel adoption, and cross-layer optimized schemes. Each of these approaches has pros and cons in terms of implementation complexity, cost, and accuracy. In the framework of CMUseum, a simple but effective method is used to accomplish the same work: we use the location information as an index to determine the transmission rate adapted in the lower layer.

In the configuration and testing phase of our system, we have measured the adapted transmission rates for all the possible locations. The rate/location tuples are provisioned to the video server beforehand. Using the location information as an index, the video server then retrieves the corresponding transmission rate and accordingly manipulates video processing. This approach largely simplifies the system complexity without sacrificing performance.

To support a large number of concurrent users, the video transcoding technique [8] is used as a function of the number of streaming videos. Based on VLC's preliminary set of transcoding options, we are working on providing more efficient and diverse transcoding functions. We are also investigating other mechanisms such as dynamic bitstream switching [9], fine-grained scalable coding, and dynamic rate control to enable faster, on-the-fly content switching.

5. ADD-ON SERVICES

The proposed infrastructure also supports various add-on services. Selected add-on services are shown in Figure 4. Using the location information, the *virtual GPS* helps each user identify his location in a labyrinth-like museum. With support of a server-side database, the *companion tracker* can provide the location information of other users, like family or friends, in a graphic manner, helping the user keep track of where their companions are.

The *walking map* helps a user find the next spot to visit. After specifying the next spot of interest, a graphic interface showing the most efficient route to that spot is displayed instantaneously. In addition, detailed directions that guide you to the destination are shown on the side of the screen. This function can also help a visitor make efficient touring plans.

The *traffic assistant* shows the snapshot of the current traffic flow in the museum. This feature enables the visitor to avoid big crowd. It also assists the curator to easily control the traffic flow.
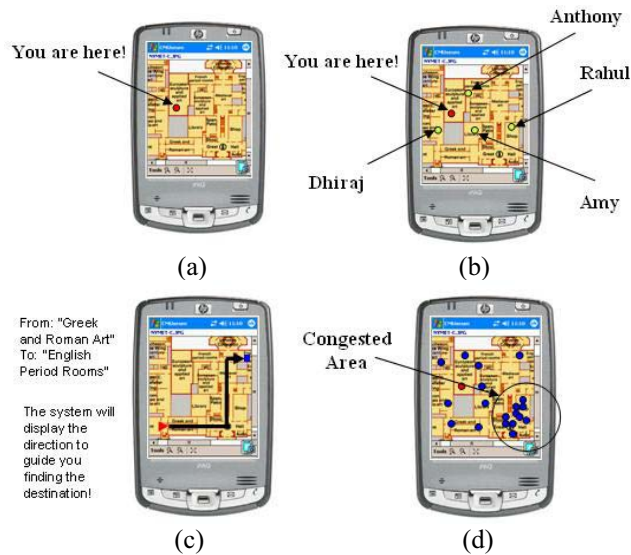
Figure 4. CMUseum's add-on services: (a) virtual GPS (b) companion tracker (c) walking map, and (d) traffic assistant

6. FUTURE WORK

Several avenues have been planned on the roadmap. In our current deployment, one access point is used in serving the video streaming work. To support a larger number of concurrent users, we plan to incorporate multiple access points with the load balancing capability [10]. Also, we will continue to investigate the effects of using the location to facilitate load balancing. The location information should enable the load balancer to do node association with optimal space reuse and better video quality.

With the current location of each user, the server side actually holds the up-to-date network topology. This topology information can be used in many aspects to enhance the networking performance. One possible direction is to enhance ad-hoc routing in multi-hop wireless LANs [11]. Wireless LANs with multi-hop extensions can provide wider coverage and higher capacity.

7. CONCLUSION

In this paper, we have presented the design and implementation of a location-aware video streaming system, CMUseum. CMUseum enables location-based video streaming applications, as well as other add-on services. CMUseum incorporates an 802.15.4 Zigbee sensor network for collecting location information and an 802.11 Wi-Fi network for streaming video contents. This framework is largely event-driven in order to support the real-time nature of the video streaming services it handles. We have demonstrated feasibility of this system through implementations of a museum touring scenario. We are now investigating how to use the location information to facilitate load balancing and ad-hoc routing to provide better quality of service.

8. ACKNOWLEDGEMENTS

We would like to thank Rahul Mangharam, Anthony Rowe, and Dhiraj Goel for their efforts on this project.

9. REFERENCES

[1] S. Banerjee, M. Youssef, R. Larsen, A. Udaya Shankar, A. Agrawala, et al, "Rover: Scalable Location-aware Computing", *IEEE Computer*, October 2002.

[2] A. Wang, C. Sørensen, S. Brede, H. Servold, and S. Gimre, "The Nidaros Framework for Development of Location-aware Applications", *In Proc. Second IFIP TC8 Working Conference on Mobile Information Systems (MOBIS),* Leeds UK, 5-6 December 2005, pp. 171-186

[3] Y. Chen, X. Y. Chen, F. Y. Rao, X. Y. Yu, Y. Li, D. Liu, "LORE: an infrastructure to support location-aware services", *IBM Journal of Research and Development*, Volume 48, Issue 5/6, Pages 601-615, 2004.

[4] "Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs) ", IEEE Std. 802.15.4, 2003.

[5] "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", *Ref. ISO/IEC 8802-11:1999(E)*, IEEE Std. 802.11-199, 1999.

[6] A. Rowe and R. Mangharam, "RT-Link: A Time-Synchronized Link Protocol for Energy Constrained Multi-hop Wireless Networks", Technical Report CMU-ECE-TR05-08, Aug 2005.

[7] VideoLAN, http://www.VideoLAN.org/

[8] J. Xin, C. W. Lin, and M. T. Sun, "Digital Video Transcoding", *Proceedings of IEEE*, vol.93, no. 1, pp. 84-97, Jan. 2005. (invited paper)

[9] B. Xie and W. Zeng, "On the Rate-Distortion Performance of Dynamic Bitstream Switching Mechanisms", *IEEE International Conference on Multimedia and Expo (ICME),* Amsterdam, the Netherlands, July 2005.

[10] Y. Bejerano, S. Han, and L. Li, "Fairness and Load Balancing in Wireless LANs Using Association Control", *ACM International Conference on Mobile Computing and Networking (MobiCom)*, Philadelphia, Sept. 2004.

[11] S. Lee, S. Banerjee, B. Bhattacharjee, "The Case for a Multi-hop Wireless Local Area Network", *Infocom 2004*, March 2004, Hong Kong, China.