

Pervasive Smart Camera Networks exploiting heterogeneous wireless Channels

Thomas Winkler
Pervasive Computing Group
Institute of Networked and Embedded Systems
Klagenfurt University
Email: thomas.winkler@uni-klu.ac.at

Bernhard Rinner
Pervasive Computing Group
Institute of Networked and Embedded Systems
Klagenfurt University
Email: bernhard.rinner@uni-klu.ac.at

Abstract—Smart cameras are embedded systems that perform on-board video content analysis and only report detected events instead of permanently streaming videos. Visual sensor networks aim at integrating smart cameras with wireless sensor networks. Camera sensors have higher requirements regarding computing power and communication bandwidth than those typically used in wireless sensor network applications. Consequently, power management is an even more critical issue. This work in progress presents an attempt to address this by combining high and low power radios as well as high and low performance computing systems in one single platform. This allows to control power consumption by selectively enabling only required components.

I. INTRODUCTION

Smart cameras are real-time embedded systems that perform image analysis on-board and typically deliver high-level event descriptions [1]. For important events, videos can be streamed allowing operators to visually evaluate the situation. Networks of smart cameras combine distributed sensing with distributed computing and have emerged thanks to a confluence of advances in the fields of computer vision, sensor networks, embedded systems and distributed computing [2]. Besides ongoing progress in areas like lenses, image sensors and processors, these networks are one key aspect of the revolution in the field of cameras taking place at the dawn of the twenty first century. This revolution will change our conception of cameras as boxes that capture images into a more general notion of cameras as spatially distributed sensors that generate data and events.

Law enforcement and security are the most obvious applications of smart camera networks. Large areas can be covered only by large numbers of cameras; analysis generally requires fusing information from several cameras. However, distributed smart cameras have many other uses as well, including machine vision, medicine, elderly care and entertainment. It is foreseeable that these networks will pervade into several new applications in the near future.

In [3] we have outlined the evolution from single towards *pervasive smart cameras (PSC)* which have similar properties as visual sensor networks [4]. Significant research is still necessary to achieve the ultimate goal of having available ubiquitous, adaptive, secure and autonomous camera networks. In this work in progress, we focus on one aspect towards this goal: providing a wireless, power-aware infrastructure

for PSCs. Wireless networking and power awareness are critical issues for visual sensor networks as they help to reduce installation and operating costs, simplify deployment and prolong operating time. We achieve power-awareness by two approaches: First, we use a dual-radio network for setting up communication among individual camera nodes. By choosing between a high-bandwidth/high-power and low-bandwidth/low-power radio we can trade communication performance for power consumption. Second, the camera nodes can be set to different power modes during operation. As a consequence, camera nodes can be powered on only when interesting activities are assumed to take place in their field of view. The overall PSC infrastructure can dynamically adjust sensing, processing and communication performance to the current requirements.

The remainder of this paper is organized as follows: Section II presents our PSC architecture which consists of smart camera nodes and a 802.11/802.15.4 dual radio communication infrastructure. We sketch the generic camera node architecture. Section III describes our preliminary results using two different scenarios. A "low-bandwidth" tracking example demonstrates some advantages of our infrastructure. Section IV discusses related work and Section V concludes the paper with an outlook on future work.

II. SYSTEM ARCHITECTURE

Figure 1 presents a PSC network architecture example. The network consists of seven cameras with adjacent or overlapping fields of view. Each camera node is equipped with a high bandwidth radio to facilitate delivery of video streams to a consumer. To keep installation and management as simple as possible, the high-bandwidth network does not rely on managed infrastructure like access points but uses a mesh topology for data transmission.

During normal operation, the wireless network is used for system management, coordination between cameras as well as delivery of detected events to monitoring stations. In this mode, the 802.11 network is not required and can be disabled to conserve power. To maintain communication links, camera nodes are equipped with an additional 802.15.4 based radio represented by grey dots in Figure 1. They have the advantage of consuming orders of magnitude less power than the 802.11

based radios [5]. Similar as for the high-bandwidth network, data between nodes of the 802.15.4 network is routed in a multihop fashion. Since achievable communication distances for 802.15.4 networks typically are lower than those of 802.11, the network is augmented with nodes only equipped with 802.15.4 radios used for packet forwarding.




Fig. 1. Proposed PSC networking architecture. Camera nodes are equipped with two radios (high and low performance) while intermediate nodes are equipped with only a single, low-performance radio.

A. Node Architecture

A PSC camera node as shown in Figure 2 consists of two distinct subsystems: A camera system and a wireless sensor node. The camera subsystem is an embedded computing platform equipped with a CMOS sensor, an embedded processor (e.g. ARM), appropriate amounts of RAM and flash memory and optional accelerators for video content analysis and video encoding. Moreover, the system provides a high-bandwidth 802.11 wireless networking interface capable of streaming videos at high resolutions. The second major building block of a PSC camera node is a wireless sensor node or mote as depicted on the right side of figure 2. The mote itself is a low-power embedded platform equipped with a low-bandwidth 802.15.4 radio. While the high-bandwidth radio is capable of transmitting at typical data rates of 54 Mbps and more, the low-bandwidth radio can achieve up to 250 kbps. Having two distinct computing platforms instead of having one platform equipped with two different radios provides a number of advantages. It opens the possibility to not only turn off the high-bandwidth radio but it also allows to put the whole camera subsystem into sleep modes controlled by the mote platform. Based on information from adjacent nodes, the image processing subsystem can be powered up or down as required. Another advantage is that the same mote platform with identical software can be used for the relay nodes.

Figure 2 also provides an overview of the software architecture of a PSC node. The camera system is running an embedded Linux OS providing a set of standard libraries used for networking, image processing and system management. The mote platform is running a mote specific operating system. A software layer manages communication between both

subsystems using a wired interconnection (e.g. USB, SPI, ...). On the mote, a number of services is offered. One such service is a *NetworkListenService* waiting for incoming connections on a given port. Services can be instantiated by the camera subsystem as required. Each service is assigned a unique id which is used to route messages from the service handlers on the mote to their counterparts on the camera subsystem. Message routing and distribution is done transparently by the *Mote Service Manager*. This allows applications on the camera subsystem to e.g. use the low-bandwidth communication channel of the mote in a similar fashion as the local high-bandwidth radio. Additionally to making mote functionality available to the camera subsystem, there also is the option to implement services that are only running on the mote such as network management including node and resource discovery. This leads to a partitioning of the workload between the mote and camera subsystem.

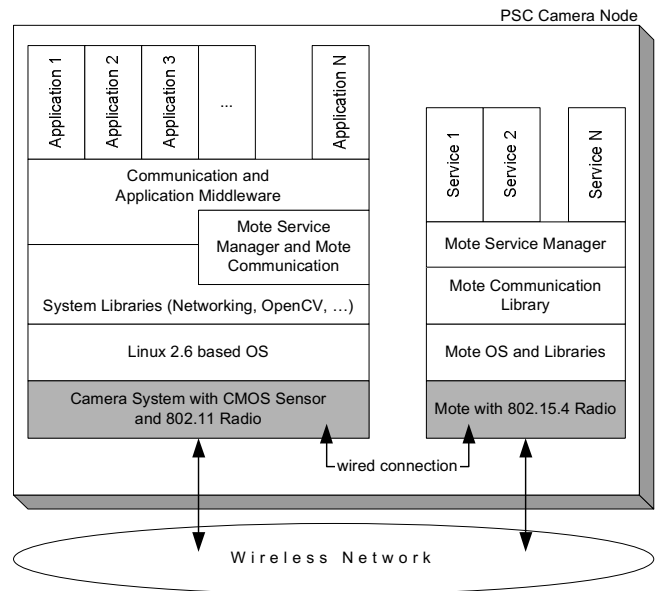


Fig. 2. A PSC node consists of two components: A camera subsystem running Linux on an embedded platform with a high-bandwidth radio and a mote equipped with a low-power radio.

III. PRELIMINARY RESULTS

This section provides information about the prototype setup used for early evaluations. To demonstrate the applicability of the proposed ideas, two application scenarios have been implemented.

A. Prototype Environment

The prototype setup consists of two laptop computers equipped with integrated WiFi 802.11g radios. Both laptops are running a recent version of Linux. Via USB, a SunSPOT wireless sensor node is attached to each laptop. The SunSPOTs are equipped with an ARM9 processor, 512 kB RAM and a Chipcon CC2420 802.15.4 radio. They provide a Java runtime environment with a Java virtual machine directly running on the ARM CPU without separate underlying OS [6]. The image analysis components running on the laptops are fed with

video sequences pre-recorded with a consumer camcorder. For computer vision related tasks the OpenCV [7] library is used.

On the SunSPOTs a service manager was implemented handling the USB communication with the laptops. For evaluation, simple broadcast sender and listener services were implemented on top of it. On the laptops, counterparts of the mote service handlers were implemented as a Python library. This allows to send and receive messages from applications on the laptops via the SunSPOTs 802.15.4 radio. To provide very rough estimates of the 802.11 and 802.15.4 radio power consumption, Table I presents power consumption figures based on values read from the internal power controllers of the devices.

	Listen	Receive	Transmit
802.15.4 (SunSPOT)	57 mW	122 mW	135 mW
802.11g (Laptop)	800 mW	1350 mW	1700 mW

TABLE I
POWER CONSUMPTION OF THE 802.11 AND 802.15.4 RADIOS BASED ON VALUES REPORTED BY THE INTERNAL POWER CONTROLLERS.

B. Scenario 1: On-Demand Video Streaming

This scenario explores the possibility of using the 802.15.4 radio to deliver events detected by the camera subsystem to a control station. If the information is considered important by the operator and requires visual inspection, the low-bandwidth radio subsequently is employed to set up an 802.11 connection. This connection is then used to deliver a live video stream of the observed area. The communication pattern is shown in Figure 3 where the two laptops of the prototype setup are denoted as node 1 and 2.

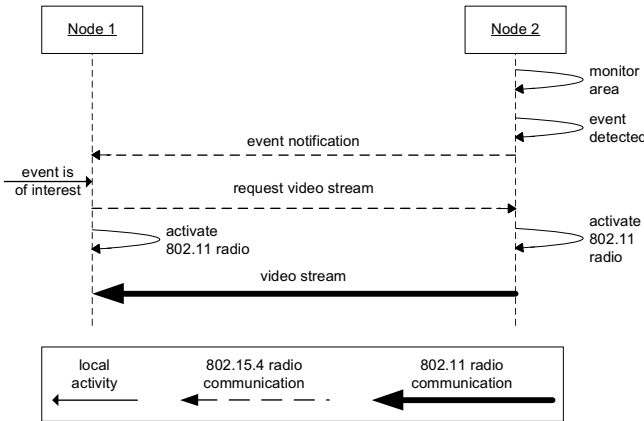


Fig. 3. Once an event is detected in the field of view of node 2, node 1 is notified via the 802.15.4 radio. Only if this event requires closer inspection, the 802.11 wireless channel is established and live video streaming is initiated.

Time required from when video streaming is requested via 802.15.4 to the point when the 802.11 link is operational and video data is streamed was measured to be around 1.5 seconds. This process involves loading the 802.11 kernel drivers, network configuration and starting the streaming server. To avoid loss of information due to setup times of the 802.11 channel, live video can be recorded on the camera and then be delivered using quick motion once the channel is operational.

C. Scenario 2: Low Bandwidth Data Delivery

This scenario evaluates the use of the 802.15.4 radio for delivery of information on motion events detected on a camera node. Once an event in the observed region is detected, a message is sent to the control station indicating the event. Subsequently, updates of the location of the detected moving object are sent to the control station via 802.15.4. There, this information is displayed as a bounding box of the object together with a simple object trace. As background, a static image of the cameras field of view is used. Example images of a person walking along a hallway are show in Figure 4.

As static background a 24bit color JPEG image at a resolution of 720x576 pixels is used. As shown in Table II, transmission time via 802.15.4 is above 12 seconds resulting in a datarate of 6.5 kB/s. The datarate achieved for the same image via 802.11 was measured as 551 kB/s.

	80 kB image	95 bytes data
802.15.4 (SunSPOT)	12.1 s	12 ms
802.11g (Laptop)	145 ms	2.5 ms

TABLE II
TRANSMISSION TIMES FOR 80 kB AND 95 BYTES OF DATA.

For transmitting location information of a tracked object, 8 bytes are required. In Table II this case is covered with the transmission time for a 95 bytes packet. This packet size was chosen as it can be sent in one frame by the SunSPOT radio stack. Transmission times of 12 ms allow to send position updates frequently enough for smooth location visualization.



Fig. 4. A sequence of images showing the location of a tracked person as seen on the control station. Location information is transmitted via 802.15.4.

Together with the power figures from Table I, the transmission times indicate that selectively using 802.11 over 802.15.4 for image transmission can reduce power consumption by a factor up to 7. It has to be noted that power consumption for initialization of wireless interfaces is currently not considered but was shown to be significant for 802.11 radios [5].

IV. RELATED WORK

This section presents representatives of work on visual sensor networks as well as research on dual-radio architectures.

A. Visual Sensor Networks

Currently the majority of work in the field of visual sensor networks is focused on the development of single sensor nodes. Representatives of such platforms are the CMUcam 3 [8] or the WiCa [9]. While the former is based on a general purpose ARM7 CPU, the latter is using a special purpose SIMD processor. A broad overview of available platforms is presented in [3]. An example of research focusing on networking aspects is Firefly Mosaic [10] where the CMUcam is combined with a Firefly mote providing an 802.15.4 radio. A similar approach is taken by Citric [11] where a custom built camera module based on a PXA270 CPU is combined with a Telos Sky mote.

A common observation is that visual sensor networks are currently focused on using single, low-bandwidth radios.

B. Heterogeneous Radio Architectures

With CoolSpots, Pering et al. [12] describe a system where devices with WiFi and Bluetooth communication channels are used for web-surfing and file-transfers. The appropriate network link is dynamically selected depending on the current network load. CoolSpots is limited to two communication partners. Energy savings are reported to be up to 75% depending on the scenario.

In [13] Stathopoulos et al. use a dual-radio platform to implement a system that selectively enables high-bandwidth radios to form end-to-end communication paths. The low-bandwidth network is used to control the high-bandwidth network. The work focuses on development of topology control and routing mechanisms for dual-radio networks.

Lymberopoulos et al. [5] examine the energy efficiency of a platform equipped with a Chipcon CC2420 802.15.4 and a 802.11b radio. They note that the startup times and startup power consumption is considerably lower for the 802.15.4 radio. They conclude that using 802.11 only amortizes for larger amounts of data. The breakeven point is determined around 725 kB. Considering that 802.11 can transmit data over longer distances than 802.15.4 without intermediate retransmissions, it is shown that this breakeven point drops to 240 kB.

V. CONCLUSION AND FUTURE WORK

This work explores the application of a dual-radio setup in the context of smart camera networks. Contrary to previous work such as Firefly or Citric where only a low-bandwidth radio is used, our approach allows for video streaming which often is required for visual verification. During normal operation where only the low-bandwidth radio is used for event delivery, power consumption is comparable to other visual sensor networks. Future work will focus on extending the ideas to a multi-camera testbed and controlling the camera subsystem via the mote platform. An area not yet addressed by research on dual-radio networks where we expect to contribute, is the

development of a software layer providing a unified abstraction for the underlying, heterogeneous networking technologies. Our next steps will concentrate on the following topics:

Multi-Camera Testbed: Currently we are working on design and deployment of a multi-camera testbed at our institute. Initially, the setup will consist of six embedded smart cameras equipped with ARM CPUs, hardware MPEG4 encoders and 802.11 radios. As mote platform, SunSPOT sensors will be used. The testbed will allow to evaluate the dual-radio architecture, including multi-hop networking, under real conditions.

Extended Evaluation Scenarios: Evaluation scenarios will be extended to take advantage of the new testbed including multi-camera tracking or office infrastructure control. Additionally, more accurate power consumption measurements of the embedded platform are planned.

PSC Middleware: Based on current prototypes, a middleware is to be designed that addresses specific needs of PSC networks. These include a common naming and addressing scheme for multi-radio environments, abstraction of heterogeneous wireless channels using a common API, support for self-configuration including services announcement and discovery, time and data synchronization as well as camera calibration.

REFERENCES

- [1] M. Bramberger, A. Doblander, A. Maier, B. Rinner, and H. Schwabach, "Distributed Embedded Smart Cameras for Surveillance Applications," *IEEE Computer*, vol. 39, no. 2, pp. 68–75, Feb. 2006.
- [2] B. Rinner and W. Wolf, "Proceedings of the IEEE: Special issue on distributed smart cameras," *Proceedings of the IEEE*, vol. 96, no. 10, Oct. 2008.
- [3] B. Rinner, T. Winkler, W. Schriebl, M. Quaritsch, and W. Wolf, "The Evolution from Single to Pervasive Smart Cameras," in *Proceedings of the Int. Conference on Distributed Smart Cameras*, 2008.
- [4] I. F. Akyildiz, T. Melodia, and K. R. Chowdhury, "Wireless Multimedia Sensor Networks: Applications and Testbeds," *Proceedings of the IEEE*, vol. 96, no. 10, 2008.
- [5] D. Lymberopoulos, N. B. Priyantha, M. Goraczko, and F. Zhao, "Towards Energy Efficient Design of Multi-radio Platforms for Wireless Sensor Networks," in *Proceedings of the International Conference on Information Processing in Sensor Networks*, 2008, pp. 257–268.
- [6] D. Simon, C. Cifuentes, D. Cleal, J. Daniels, and D. White, "Java on the Bare Metal of Wireless Sensor Devices," in *Proceedings of the International Conference on Virtual Execution Environments*, 2006.
- [7] G. R. Bradski and A. Kaehler, *Learning OpenCV - Computer Vision with the OpenCV Library*. O'Reilly, 2008.
- [8] A. Rowe, A. G. Goode, D. Goel, and I. Nourbakhsh, "CMUcam3: An Open Programmable Embedded Vision Sensor," Robotics Institute, Carnegie Mellon University, Tech. Rep., 2007.
- [9] R. Kleihorst, A. Abbo, B. Schueler, and A. Danilin, "Camera Mote with a High-Performance Parallel Processor for Real-Time Frame-Based Video Processing," in *Proceedings of the First International Conference on Distributed Smart Cameras*, 2007, pp. 109–116.
- [10] A. Rowe, D. Goel, and R. Rajkumar, "FireFly Mosaic: A Vision-Enabled Wireless Sensor Networking System," in *Proceedings of the International Real-Time Systems Symposium*, 2007, pp. 459–468.
- [11] P. W.-C. Chen, P. Ahammad, C. Boyer, S.-I. Huang, L. Lin, E. J. Lobaton, M. L. Meingast, S. Oh, S. Wang, P. Yan, A. Yang, C. Yeo, L.-C. Chang, D. Tygar, and S. S. Sastry, "CITRIC: A Low-Bandwidth Wireless Camera Network Platform," in *Proceedings of the International Conference on Distributed Smart Cameras*, 2008.
- [12] T. Pering, Y. Agarwal, R. Gupta, and R. Want, "CoolSpots: Reducing the Power Consumption of wireless mobile Devices with multiple Radio Interfaces," in *Proceedings of the International Conference on Mobile Systems, Applications and Services*, 2006, pp. 220–232.
- [13] T. Stathopoulos, M. Lukac, D. McIntire, J. Heidemann, D. Estrin, and W. Kaiser, "End-to-End Routing for Dual-Radio Sensor Networks," in *Proceedings of the Int. Conference on Computer Communications*, 2007.