

UC Berkeley

Working Papers

Title

Berkeley Highway Lab Video Data Collection System

Permalink

<https://escholarship.org/uc/item/4t44v5rm>

Journal

California Partners for Advanced Transit and Highways (PATH)

Authors

Chen, Chao

Lyddy, Daniel

Dundar, Baris

Publication Date

2004-11-01

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Berkeley Highway Lab Video Data Collection System

Chao Chen
Daniel Lyddy
Baris Dunder

California PATH Working Paper
UCB-ITS-PWP-2004-8

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department Transportation, Federal Highway Administration.

The contents of this report reflect the views of the authors who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the State of California. This report does not constitute a standard, specification, or regulation.

Final Report for Task Order 4129

November 2004

ISSN 1055-1417

Berkeley Highway Lab Video Data Collection System

PATH Task Order 4129 Final Report

Chao Chen, Daniel Lyddy, Baris Dunder

Table of contents

Berkeley Highway Lab Video Data Collection System.....	0
Table of contents.....	1
Executive Summary	2
Background.....	3
Pacific Park Plaza.....	3
Original BHL Video Collection System	4
Shortcomings of the Original System	4
New System Design.....	5
Motivation	5
System design	5
Camera choice.....	6
Servers	9
Mounting structure.....	9
Coverage.....	10
Physical layout.....	11
Software development.....	12
Communications	14
Results	14
Software license problem	14
Problem related to firewire repeater.....	14
Communication problems	14
Zoom and aim	15
Image quality	15
Future plans	15

Executive Summary

The goal of this project is to replace the existing analog video collection system on the roof of Pacific Park Plaza (PPP) with a digital one. This video collection system is part of the Berkeley Highway Laboratory (BHL) testbed. It records video of traffic on a continuous one-kilometer section of Interstate 80 (I-80) near Emeryville, CA. This section of I-80 features on-ramps, off-ramps, and weaving zones. The video of interactions between vehicles can be fed to a machine-vision system, which generates vehicle trajectories to be used in a variety of traffic studies. The video itself can also be used for both human and machine-vision based verification of loop data.

The previous system used 12 analog video cameras connected to SVHS VCR decks. This system operated between July 2000 and August 2002. During this period, we collected approximately 180 days worth of tapes, spanning approximately 900 hours from each camera, for a total of more than 10,000 camera-hours. This system covered approximately 1.5 km of I-80 in both directions, but because of mounting requirements, there was a 50-meter gap in coverage right near the side of the building.

We had several motives for replacing the old analog system with an all-digital system. First, the new digital data collection process is much more streamlined than the old analog system. Analog video must be digitized before it can be used in a machine vision system. This analog-to-digital conversion process has proven to be slow, labor-intensive, and error-prone. With the Digital Video System, video data is written directly to disk drives right on the PPP roof. From there, it can be offloaded to a remote machine automatically.

Second, with the old analog system, someone had to retrieve the videotapes from the roof of PPP and deliver them to the RFS. Liability concerns prevented us from using University Employees or Students, so we had to pay a contractor to perform this task. Of course, as part of the contract, we had to cover their insurance costs, which added significantly to our operating budget.

Finally, after all this collection and digitizing, the final video product proved to be unsuitable for use in certain types of machine vision tasks. One of our vehicle detection algorithms relied on the ability to find "edges" in a given still image. The video from our analog traffic cameras was interlaced, which means that even and odd scan lines from the same video still are actually captured 1/60th of a second apart in real time. In this time period, a vehicle traveling at 60mph moves about 1.5 feet, which depending on camera resolution, could result in several pixels of "zigzag" in what should have been recorded as an edge. The new digital cameras use progressive scan technology, which means all scan lines are captured at the same moment in time. Progressive scan not only eliminates the zigzag problem, it also produces better images in terms of overall quality. In summary, the new digital system saves time, effort, and money over the previous tape collection and manual digitizing process, and the resulting video quality is superior.

The new system became operational on July 29, 2004. There are eight digital cameras mounted on the roof of PPP. These eight cameras are connected through a fiber-optic IEEE1394 "Firewire" repeater system to two servers in the building, which are in charge of collecting the video locally and storing it to disk. These two local servers are networked through a wireless link to a remote server at the Richmond Field Station (RFS), which itself has connections to the entire Berkeley network. For extra security, the PPP servers, wireless radios, and RFS server all sit on a private network that is firewalled from the rest of the network. This allows us to restrict access to the PPP servers to authorized users and machines.

The eight digital cameras cover about one kilometer of I-80 in both directions, but because of the new mount design, the coverage is completely continuous with no gaps. The new system allows us to store digital data on the PPP servers, from which we can use File Transfer Protocol (FTP) to move the files to machines on campus. So far, we have performed one test and successfully recorded about three hours of video.

This document describes how the new system works and explains the rationale behind some of the design decisions that we made. We still need to make many improvements to the system to make it user friendly and fully functional. We outline the planned work here.

Background

Pacific Park Plaza



Figure 1 Pacific Park Plaza, Emeryville, CA.

The Pacific Park Plaza (PPP) is a 30-story condominium located at 6363 Christie Avenue in Emeryville, California. The building itself is approximately 100 meters tall, and it sits about 10 meters from the edge of a heavily traveled section of Interstate 80. The building's height and proximity to I-80 make it the perfect vantage point from which to observe and study traffic on this interesting section of highway.

Original BHL Video Collection System



Figure 2 Original camera mount.

In 1999, PATH installed the first iteration of the BHL video collection system on the roof of PPP. The system used 12 fixed-alignment analog S-Video cameras plus two Pan-Tilt-Zoom (PTZ) cameras mounted on four poles installed at the westernmost point of the building's roof. Power, video, and control signals were carried by bundled cables running to a supply room at the center of the roof, about 200 feet away from the cameras. The two video signals from each camera were carried by a pair of coaxial cables in the bundle that were each roughly equivalent to RG-11 video cable. The server room housed a cabinet with two columns of 19" racks. These racks held 12 SVHS VCRs, a video multiplexer, two video/power/control breakout boxes, and a Macintosh computer to control the entire system. This Macintosh computer was connected to a Linux server at the RFS over a 5 Megabit per second (Mbps) wireless Ethernet link. This connection was set up so that a user at PATH or Caltrans DRI with appropriate permissions could control and view the output from the two PTZ cameras through a web interface. PATH Researchers could also control the VCR taping schedule either manually through the web, or automatically through a scripting language developed by Radar-Digital Systems of Soquel, California.

Shortcomings of the Original System

For its time, our system was state-of-the-art. Even so, it suffered from several shortcomings. The most serious of these shortcomings was the data collection process, which was cumbersome. Once the video information was recorded to tape, those tapes had to be removed and replaced with blank ones. The recorded tapes were then transported back to the RFS, where the analog video was digitized for use in our machine vision algorithms. Digitization is already a time consuming and error-prone process for one tape, and for each day of video collection, we would have to process twelve tapes.

The second major shortcoming of the original installation was the camera mounting system. The roof of PPP is enclosed by a four-foot parapet. For the "close-looking" cameras to be able to "see" over this parapet, they had to be mounted either as far out from the building as safely possible, or as high as safely possible. At the time of the first installation, PPP Management was not willing to allow us to mount cameras out away from their building, so instead we chose to mount the two closest-looking cameras near the tops of two ten-foot mounting poles. We mount the two remote PTZ cameras on near the tops of two other mounting poles, to give them the smallest possible "blind spots" when looking over the roof's edge. We distributed the remaining ten fixed-alignment cameras among the four poles so that there would be three such cameras per pole. The photo above shows one of the camera poles that holds three far-looking fixed-

alignment cameras mounted low and a remote PTZ camera mounted high. We had one other camera pole with this configuration. The other two camera poles had three close-looking cameras mounted high.

To support these heavy camera poles, we used the same davit posts that the building's window washers use to support their platforms. Roughly four times per year, PPP management would have their windows washed, which meant that our contractor would have to take down the cameras so that the window washers could use the davit posts that we had been occupying. Once the window washers finished, we would return to the roof with the contractor to re-install and re-align the fixed cameras. This re-alignment required the contractor to stand on a ladder and physically aim the camera while a researcher watched a monitor to direct the camera's movement. This process was difficult, potentially dangerous, and essentially non-repeatable.

New System Design

Motivation

As stated above, our old design used ten-foot vertical poles that had camera support arms mounted at right angles. The cameras themselves were attached to the arms. The pole is about 10 feet high, and the arms are about 2 feet. The cameras are mounted inside Cohu environmental enclosures which kept out dust and moisture. They are wired with power and coaxial cables.

The new system starts with eight cameras and can be expanded to include several more. It uses digital firewire cameras which need only a six-pin firewire to carry both the video and power. The cameras also sit inside environmental enclosures, made by Pelco. They are smaller than the Cohus, which means the mounting structure can be smaller, and all cameras can fit on one mounting structure. Because digital cameras have higher resolution than analog ones, the same length of roadway can be covered with fewer cameras to achieve the same resolution.

Table 1 Comparison between old and new systems.

<i>Property</i>	<i>Analog</i>	<i>Digital</i>
Camera	NTSC (400 lines)	640x480
Progressive scan	No	Yes
Enclosure	21 inches	12 inches
Weight of camera + enclosure	25 lb	10 lb
Number of mounts	4	1
Height of mount	10 feet	54 inches
Road coverage	1.5 km with break	1 km continuous
Data collection	36 hours per month	limited only by disk space
Digitization	laborious	automatic

System design

Among the choices we had to make when designing this system were what kind of cameras to get, what kind and how many machines we need at PPP, and what communication method to use to transmit the video. We chose Sony DFW-VL500 firewire cameras with integrated zoom, and use two Dell dual Xeon servers as the PPP video servers. The communication between PPP and CCIT will be a 20Mbps wireless link.




Figure 3 System design diagram.

Figure 3 shows the system. Because the distance between the PPP server room and the cameras is much greater than the range of firewire, we used four firewire-to-fiber repeaters to bridge the gap. These choices are explained below.

Camera choice

There are four major segments in the video camera market: Consumer, Professional Broadcast, Surveillance, and Industrial Machine Vision. Cameras are designed for the Consumer segment with the expectation that a novice will be constantly interacting with the camera while it is recording, therefore user-friendliness is of prime importance in this segment. In the Professional segment, designers expect a higher level of expertise from their users, and the users expect a higher level of video quality from their cameras. However, in both of these segments, remote functionality is usually limited to file transfer in Consumer-grade cameras, plus some in-camera editing ability in Professional-grade cameras. In both of these segments, access to basic settings such as zoom, iris, and focus is restricted to buttons or sliders on the cameras themselves.

Surveillance cameras are designed to be controlled remotely in indoor or outdoor environments that can be harsh. For most surveillance applications, the video quality is not of paramount importance, because these cameras are generally only used to generate a first warning that something suspicious or out of the ordinary is happening. Usually this first warning results in humans being dispatched to the scene to investigate further.

On the other hand, Industrial Machine Vision cameras are used as sensors in automated processes such as parts sorting and inspection. Users of these applications want the camera to do as much of the "dirty work" as possible, so video quality is the first priority in these applications. Unfortunately for us, these cameras

are generally designed to work in a controlled environment such as a clean room, where temperature, air quality, and lighting can be strictly controlled.

To keep costs down, camera vendors generally target a single segment, without much overlap. Therefore, cameras that produce high-quality progressive-scan video feeds for Machine Vision applications generally do not provide the remote lens controllability and environmental housing typical in Surveillance cameras, and vice-versa.

In the original installation, we started with standard Cohu 8250 analog traffic cameras, but instead of using the standard single-channel composite NTSC feed, we had the cameras modified to produce two-channel S-Video (also known as Y/C). The S-Video scheme feeds the baseband luminance (grayscale) and frequency-shifted chrominance (color) signals through separate coaxials, which reduces crosstalk, and improves overall image quality.

Unfortunately, like most surveillance cameras, our Cohu cameras did not produce progressive-scan video, so that even and odd fields from the same image frame were actually captured 1/60th of a second apart. This is not a problem for human viewers because the human vision system naturally "averages out" the interlace between the two fields. However, the motion between fields degrades the performance of a machine vision system, especially one that is attempting to detect and track the motion of "rigid" bodies.

For the latest installation, we searched for a model of camera that had an embedded, controllable lens like a Surveillance camera, but which produced progressive-scan video like a Machine Vision camera. We found two such cameras on the market, but only one was available for purchase at the time of installation. This camera, the Sony DFW-VL500, connects to an external controller via an IEEE-1394 "Firewire" port. The 6-pin firewire connection carries power and control signals into the camera from the host controller, and the same connection sends digital video back to the host. The cameras also have connections for an external trigger input, which would in principle allow us to synchronize multiple cameras to an external timebase. We determined that the benefit of this capability did not justify the extra hardware and software effort and cost required to take advantage of it.



Figure 4 Our choice: Sony DFW-VL500 camera.

Firewire was originally designed as a high-speed "desktop" standard, used to connect peripherals to a host machine presumably in the same room. The Firewire standard stipulates that any connection longer than 4.7 meters or 15 feet requires intervening repeaters. Because of the way the PPP roof is physically laid out, we were forced to install our host controllers in a supply penthouse that is about more than 200 feet away from the camera mounts.

Rather than attempting to use more than ten Firewire-to-Firewire repeaters in a chain, we instead used a Firewire-over-fiber communications link manufactured by a company called Opticis. This link extends the range of Firewire by converting the electrical signals to optical ones and transporting them over fiber for the bulk of the distance. The only disadvantage of the system is that we must supply DC voltage at the "far end"

to power the conversion from optical back to electrical, and to power the cameras themselves. Aside from this inconvenience, and a slight increase in the propagation delay times, this system is transparent from the perspective of the host controllers and cameras. As far as those devices are concerned, the Firewire-over-fiber link behaves just as any other "normal" Firewire connection would.

A more major disadvantage arises from the bandwidth used by each camera. At its highest image quality setting, each camera provides 30 frames per second at a resolution of 640x480 pixels per frame, using YUV 4:2:2 color encoding. This means that the "Y" or grayscale information is transmitted at the full 640x480 resolution, while the "U" and "V" color channels are down-sampled by a factor of two after anti-aliasing filters are applied. This produces an effective color depth of 16 bits per pixel. Thus, the raw transmission rate (without framing or encapsulation overhead) is 150 Megabits per second (Mbps). Firewire bandwidth is restricted to 480 Mbps per channel, and because of the overhead, we were only able to connect two cameras in parallel to each Firewire channel.

Once these high-bandwidth feeds arrive inside the host controller, there is still the question of how to store them. At these data rates, each camera would consume just under 20 Megabytes per second (MB/s) of disk space, or roughly 70 Gigabytes per hour (GB/hr). This means that each hour of storage for all eight cameras would cost us 560 GB, so that a 1.0 Terabyte disk array would not even be enough to store two hours' worth of video. To extend the storage time, then, we must compress the video using some kind of coder/decoder (codec). Each codec is a software module that uses the host controller's CPU to compress each video feed "on-the-fly," and compressing eight such video feeds in real-time creates a substantial load for the controller. Since we have two such host controllers, we were able to divide this load in half for each controller, but even then, this extra load decreases the amount of processing power available for other tasks such as writing to disk or transmitting over our communications link.

Because of these shortcomings of the camera, we also evaluated alternatives. One promising camera was the Lumenera Le175, which claims to compress each video stream using hardware codecs embedded in each camera before transmitting them over standard Category-5 ethernet cable. If these cameras worked as advertised, they would have solved both our range and data rate problems. Unfortunately, because the design of these cameras is so new, Lumenera was not able to deliver prototypes in time for us to evaluate this camera for use in our project. However, the marketing manager at Lumenera indicated that they were also working on an "all-in-one" Pan-Tilt-Zoom solution that could be completely controlled over an ethernet link. This camera would produce digital progressive-scan in different combinations of base resolutions, color depths, and compression ratios. We intend to follow the development of this unit and possibly purchase one or two of these cameras to use as substitutes for the Pan-Tilt-Zoom cameras that were present in the first installation. If they can be ruggedized for continuous outdoor use, these all-in-one PTZ units might prove to be a cost-effective, superior-quality alternative to the "analog camera plus hardware digitizer and codec" scheme currently used in Caltrans traffic surveillance applications.




Figure 5 Lumenera camera, an alternative that we also considered.

Servers

The cameras are connected to two video servers in a supply room on the roof of PPP. These servers control the cameras and record the video. They are connected wirelessly to a computer at Richmond Field Station, which is also on the University network. This machine is called `highwaylab.berkeley.edu`. We access the video servers through the highwaylab server.

Because raw video storage requires a lot of disk space, we need to compress the video at the video servers on PPP before transmitting it over the wireless link. This task is CPU-intensive. We are using two Dell dual Xeon servers, each controlling four cameras.

We had a choice of using Windows, Linux, or Macintosh machines. We chose Windows mainly because of availability of drivers. The machine vision cameras that we use generally come with only Windows drivers. The video servers run Windows Server 2003. We use Remote Desktop to control them from our computers in the office.



Figure 6 Video servers.

Mounting structure

To eliminate these problems, we have installed a single fixed mount near the westernmost point of the building, in a location that is unaffected by the window washing cycle. Since the first installation, PPP Management has allowed a local television station to mount a live traffic camera that is cantilevered out beyond the building. PPP Management has become less "suspicious" of such installations, owing to the fact that neither the TV station camera nor any of the fourteen cameras that we initially installed have caused any problems to date. Our new mount design takes advantage of this confidence by placing cameras out over the parapet, where they can be supported by that structure. A by-product of this is that the cameras are now accessible and aimable by someone standing on the roof itself (instead of a ladder). This system is safer, more robust, and more repeatable than the old mounting scheme.

The cameras are mounted in a different location on the roof from the previous system. The analog cameras were mounted six each on two davit poles. The davits are also used for window washing. In the beginning, we envisioned only to use them for six months, so window washing wasn't a big concern in the design. But they have been up for several years, and must be taken down and put back up for each window washing,

which occurred quarterly. The new location is not used by window washers, so we don't need to move them as much.




Figure 7 New camera mount structure and complete system.

The new mounting structure is also different from the old one. It has two arms made of 12-inch steel channels, where the cameras are mounted. The channels rest a few inches on top of the parapet, and are supported by rubber feet. It is more rigid than the pole structure. The cameras themselves are inside Pelco tube-shaped enclosures and attached to enclosure mounts of varying heights. Each arm of the mounting structure can be rotated toward the middle when the cameras are serviced, where it's easier to work on them than over the parapet.

When we designed the mount, there was already a TV camera at the same davit used by KGO station. It has an arm that overhangs the edge of building, and a big camera mounted on a pan-tilt head. One of the requirements of our mount is that it must allow KGO access to their camera. This is achieved by the swiveling of the right mount arm. When the right BHL plate is rotated toward the center of the building, the KGO camera arm can swing over it when the camera needs to be serviced.

Coverage

The section of I-80 near PPP is very interesting in terms of traffic characteristics. About a half mile to the south is the 80/580/880 interchange, a location with a lot of merging. The left lanes in both directions become HOV during certain times of day, which induces lane changes and influences traffic during the transition times. There are several exits and on-ramps at Powell Street and Ashby Avenue. About a mile to the north is University Avenue. But the angle makes it difficult to track vehicles that far.

While the 12 analog cameras covered between the 80/580/880 interchange and University Avenue, at the far ends the video was not good enough for processing. The effective range was approximately between Powell and Ashby. This is also the initial coverage of the eight digital cameras. We planned the coverage with the goal of capturing a continuous section of the freeway in both directions, such that a computer program can track vehicles across the entire section. This places some constraints on the angle and zoom of the cameras, namely, neighboring cameras must have some overlap, and a minimum resolution is about 1 pixel per foot.




Figure 8 North-looking camera views, with incorrect zoom.

Since each camera has 640 pixels across, this minimum resolution gives about a 200 meter maximum range for each camera, or 1600 meters for eight cameras. However, because of the angles and overlap requirements, we cannot achieve this. On the other hand, the distance between Powell Street and Ashby Avenue is roughly 1000 meters, which is achievable with the eight cameras.

Figure 8 shows the four north-looking cameras. The zoom can be adjusted remotely but the aim cannot. Since we aimed the cameras, a contractor performed additional installations and moved several of the cameras. We need to re-aim them. This process requires one to use a laptop at the cameras or two people, one at the cameras and the other at the servers. It's not very exact because you have to judge the overlap and resolutions visually. It would be better to have a more reproducible process.

Physical layout

The roof of Pacific Plaza has three corners. Our cameras are mounted in the corner closest to the freeway; it's about 10 meters away from I-80. The servers are in a supply room in the center of the roof, roughly about 200 feet away. We ran a 200-foot conduit between the server room and the cameras, which houses four fiber optic cables for the four firewire-fiber repeaters. The conduit terminates in a junction box near the cameras, as does an outdoor power cable from the supply room. The junction box houses the four repeaters and their power supplies and two power strips. The power strips supply AC power to the camera enclosures, which are heated to prevent condensation.

Inside the supply room, the servers are mounted in a cabinet with standard 19" racks. The VCRs used to sit on shelves in the rack but were removed during the current installation. The original monitor, keyboard, and mouse are still used.

Software development

Currently, we access the cameras through the Windows Remote Desktop Connection to the two servers at PPP. We can control the cameras with an application that came with the cameras called the Unibrain Fire-i Application, from the company Unibrain which makes software for firewire cameras. This application shows the cameras that are connected to the machine and allows the user to change their settings such as zoom, focus, color, shutter speed, etc. We can also use this application to record video or capture frames. When the videos are recorded, we can download them to a machine in our office for archiving and processing.




Figure 9 Unibrain software interface.

While the standard application is enough to do what we want - record video from all cameras - we developed a customized application to streamline the process. It uses the Unibrain Application Programming Interface (API) for firewire cameras. Our application allows the user to record from all cameras at a specified date and time, for a specified length (15 min to 6 hours), FTP the videos to a server, and change the camera settings from the same interface. It minimizes the number of steps required to record video.




Figure 10 Custom application for quickly recording video.




Figure 11 Interface for adjusting camera properties.

Communications

There is an existing wireless data link between boxster.path.berkeley.edu at RFS and the servers at PPP. It is 5Mbps split between up and down links; the average data rate is about 400KB/s for file transfers. We are also in contract negotiations with Aircloud Communications company who will provide a 20Mbps link, which guarantees an FTP rate of over 1MB/s. Initial tests on compression suggests that the compressed video has a data rate of less than 10MB/minute, or less than 80MB/minute for eight cameras. Based on these conservative estimates, we will be able to download video in real time or near real time.

Results

This system has been installed since early August. However, it is not yet functioning properly. While we can access the two servers and see the images, as well as record video, we have some communication problems. We also need to find the best zoom/aim settings and camera settings like focus, shutter speed, etc.

Software license problem

We tested the two servers in the lab for one week before they were deployed on PPP. They worked well in the lab but after they were installed, one of them could not read the hardware USB license key. Therefore, it was operating as a 30-minute demo and disconnected from the cameras after 30 minutes. The machine had to be rebooted before the cameras are accessible again. We emailed the technical support team of Unibrain for about a week to try to solve the problem. We had several options. One was to uninstall and reinstall everything; the other was to use the newer version of the software. We are currently using version 2.5, which is no longer being supported. A new version, the 3.0, is in beta stage and has some bugs. We tried it and couldn't get it to work completely and the tech support did not solve the problems with it.

About a week after the first installation, we revisited PPP and reinstalled everything on the server that was not recognizing the license key. After several rounds of this, it worked. Now we don't have the license problem anymore.

Problem related to firewire repeater

In most applications, the cameras would be connected directly to the computer. When the computer reboots, the cameras are reset as well. But we have a repeater between the computer and the cameras, which means that the cameras are never rebooted. This caused some problems. When the machine is rebooted, sometimes one or more of the cameras do not communicate properly with the machine. Specifically, one cannot change certain camera settings like focus and zoom after a reboot. If we reset the repeater, the problem goes away. But the problem was that we are not at PPP all the time to reset the repeater.

We solved this problem by using an X10 Firecracker home automation kit that lets you control a relay from a computer. This system consists of a wireless transmitter plugged into the serial port of the computer, a wireless receiver that is plugged into the 110-V outlet, and a number of relays called Appliance Modules. The transmitter sends a signal to the receiver, which communicates with each AM by its address over the building's electrical wiring. The system is controlled by software that runs on one of the PPP servers, and we have one AM that is connected to both repeaters.

Communication problems

On the first day the cameras were installed, we downloaded several hours of video from four cameras (the other four had the license problem). The video was compressed using the DivX codec. Several weeks later, however, we started having communication problems. When downloading large files over FTP from PPP to boxster, the FTP connection would die after about 60MB. We suspect this is due to problems with the wireless link. This is an on-going problem.

We hope that this problem will be solved once the new wireless link is up. The existing link is 5 Mbps. We own the radios and are responsible for maintaining them. The link has worked fine for many years and we are not sure if the current problem is in the wireless link or something else. But it would be hard for us to diagnose the wireless link. The new link will be a leased line with a guarantee of 20Mbps and minimum down time. It will be operated by the Aircloud company and costs \$325 per month over two years. This model suits our project because we won't have to be responsible for maintaining a wireless link. Aircloud already provides wireless data access to the building and has ready access to it. The cost is a fraction of an equivalent wired link of the same data rate.

Zoom and aim

It's difficult to know when the cameras are at an optimal aim and zoom combination. We would like maximum coverage while maintaining a minimum resolution. But right now the only way we have to do this is by visually checking the images as we vary the camera orientation and zoom. Since there are four degrees of freedom, it's hard to find the optimal setting. The problem is exacerbated by that it's difficult to see the screen on the roof. We achieved the current aim by having one person manipulating the cameras and another person in the server room observing the images. The person in the server room used a walkie-talkie to tell the person at the cameras how to move them. It was a tedious process. In the future, we play to use a laptop at the cameras so the same person can move the cameras and view the images. We will get a hood to put over the laptop screen so it's more visible.

Image quality

A big reason of going from analog to digital is for the better image quality. Indeed the progressive scanning and higher resolution gives us better coverage of traffic scenes. However, the images at automatic settings appear pixilated. We think this is because of the "sharpness" setting on the cameras, which seems to artificially increase the contrast near vertical edges in the image. We are currently playing with the cameras to find the best sharpness setting.

Future plans

The BHL video data collection system is valuable because it provides ground truth data that can be used to evaluate transportation technologies and theories. The cameras' location allows us to collect high quality data that can be easily analyzed by machine vision algorithms.

Already, we are planning to work with a PATH research project that is evaluating an automatic speed enforcement technology. This project is using radars and cameras to measure the speed of vehicles on the highway and taking pictures of those that exceed the speed limit as proof. The leading researcher, Jim Misener, will use the BHL section of I-80 to conduct ground-truth tests. He will use video from PPP to verify the speeds of vehicles.

We want more researchers to use the BHL video collection system. To this end, we are building an integrated website that includes the BHL video and loop system, as well as UC Irvine's Automatic Traffic Management System (ATMS) testbed in Orange County, CA. Researchers and practitioners will use it to find testbed facilities that can help them conduct field tests.

The BHL video collection system will also be used as a part of the Next Generation Simulation (NGSIM) project funded by the Federal Highway Administration (FHWA). The goal of the project is to use real vehicle trajectories to develop microsimulation traffic models. The BHL team has already used technology developed at PATH to extract vehicle trajectories from a half hour of video taken from PPP. These videos were taken with six temporary cameras before the current system was installed. We demonstrated the extraction of 100% of vehicles over a half hour and one kilometer, a feat never achieved before. We are currently doing another project that will streamline the extraction process, such that we will have a production system able to extract many hours of video from BHL and other test sites.

In the PPP camera system, we have a valuable research and data collection tool. The ability to collect real time video on demand for such a large coverage area is unique as far as we know. We hope that this platform provides the video data for greater understanding of traffic phenomena and the application of machine vision technologies.