Loughborough University

This item was submitted to Loughborough's Institutional Repository (https://dspace.lboro.ac.uk/) by the author and is made available under the following Creative Commons Licence conditions.

For the full text of this licence, please go to:
http://creativecommons.org/licenses/by-nc-nd/2.5/

A Framework for Cross-Layer Measurements in Wireless Networks

Konstantinos G. Kyriakopoulos, William G. Whittow, David J. Parish

Department of Electronic and Electrical Engineering

Loughborough University, Loughborough, LE11 3TU, U.K.

e-mail: {elkk, elwgw, d.j.parish}@lboro.ac.uk.

Abstract

This paper formulates a framework for wireless network performance measurements with the scope of being as generic as possible. The methodology utilises a cross-layer approach in order to address the limitations of traditional layered techniques. A lot of work in the research community uses the channel power (Cp) to predict performance metrics in higher layers. There are currently two methods to measure Cp; either by using a spectrum analyser or from WiFi card information (RSSI). The paper discusses the correct configuration of a spectrum analyser (SA), to measure Cp. This paper, also provides a comparison of both SA and RSSI results produced inside an anechoic chamber for three different applications. The behaviour of the RSSI values showed significant discrepancy with both the SA results and what was intuitively expected. The results pinpoint the necessity of a cross-layer approach and the importance of carefully selected and positioned equipment for the accuracy of the measurements.

1 Introduction

In the last decade wireless communications have become omnipresent and the trends indicate that different wireless communication technologies will converge and work in conjunction with a wired network backbone. There has been a lot of research on how to improve the already established protocols (TCP/UDP) and to optimise them for a wireless channel medium. However, the underlying issues of the wireless medium and how these influence higher layers is not yet understood in the research community.

In order to improve the performance of hybrid wireless and wired networks the research community needs to examine the nature of the protocol stack as a whole and not separate it into layers. Examining just the higher layers does not reveal much information about the issues of the lower layers and similarly examining only lower layers and their parameters does not provide any information on their effect on the application layer. Thus, to understand the impact of the wireless channel medium on the upper layers, a cross-layer approach is required.

This paper formulates a framework for wireless network performance measurements in a cross-layer manner with the aim of being as generic as possible, but applied here to a specific wireless technology. In addition, the paper provides a discussion of different methods of assessing the physical layer performance of wireless networks and provides a comparison of their results after practical implementation. The paper proposes the use and discusses the correct configuration of a spectrum analyser (SA) for measuring channel power under experiments taking place inside an anechoic chamber.

The rest of the paper is structured as follows: In Section 2, related work on current methods for measuring wireless network performance is presented along with some discussion on their shortcomings. Section 3 is divided into two parts: the first part describes the test-bed in which the experiments took place in and the second part explains the gate mode and how to carefully set the spectrum analysers parameters. Section 4 presents the practical results of the cross-layer measurements from three applications: ping, iPerf and iTunes (video streaming). Finally, Section 5 presents the conclusions.

2 Related Work

Initially, the research community tried to measure wireless networks from a wired vantage point, i.e from a wired host having some knowledge of the wireless network and in some cases through the use of SNMP logs [4]. However, SNMP logs provide summarised information polled periodically and does not expose the instantaneous characteristics of the wireless channel, which is important for traffic characterisation and network diagnosis (security) purposes.

The research developed by Maryland University [7, 8] suggests deploying multiple wireless network hosts capturing traffic from the wireless network. The advantages of this methodology are that sniffers do not interact with the network, as they are just passive monitoring devices. In

237

addition, they provide detailed, per frame PHY/MAC information of the data such as Received Signal Strength Indicator (RSSI), noise and signal strength, throughput and error rates. The above papers try to address the following challenges that arise from this methodology: i) Find an ideal location for monitoring the wireless traffic, ii) address the limited capabilities for each sniffer, i.e. signal strength, processing power, disk space, and iii) combine the captured traces from various sniffers to provide a better view of the wireless traffic as some sniffers have a better view of an area than others.

The above references do not provide a cross-layer approach to wireless network measurements. The work in [6] is very close in nature to the above references but follows a cross-layer approach. In [6], two models for mapping Signal to Noise (SNR) to throughput are presented and verified with measurements from three applications: i) iPerf, ii) FTP and iii) LANFielder. The same models could be extended to new applications.

However, [6] has the following disadvantage; according to the IEEE 802.11 standard, the RSSI value is a 1-byte value (max value 255) that maps the RF energy received by the chipset of the wireless card. It is intended for use by the Wi-Fi card internally between the link and physical layers. It is not intended to be of any particular accuracy for measuring actual signal strength. Furthermore, the 802.11 standard does not define any particular mechanism nor requires that all 255 values should be used. Thus, each vendor has different maximum RSSI values and usually uses its own (often undisclosed) methodology to map the RSSI to the received power (in dBm) [5]. The RSSI values are therefore a source of confusion and uncertainty [5].

Finally, Naples University [1, 2, 3] follows a different approach with the aim of correlating the values of major physical layer quantities (i.e. Channel power, SNR, Signal to Interference Ratio (SIR)) in the wireless channel to those characterising the key higher layers parameters. The goal was to assess the performance of one protocol layer as a function of that of another or several other layers. The uniqueness of this approach is that it does not depend on the RSSI values taken from the firmware of the NIC but on signal and noise measurements taken from a spectrum analyser, independently in a semi-anechoic chamber.

3 Methodology

3.1 Testbed

It is difficult to measure electromagnetic waves from a single source as there is interference from many external objects, such as mobile phones and local area networks etc. It is also difficult to control the reflections from surfaces such as tables and walls. An anechoic chamber is a solution to these problems and consists of a metallic box with its interior walls covered in radar absorbing material (RAM). The metal box shields the experiment from external sources, while the RAM absorbs a very high percentage of the waves from the source and effectively places the experiment in an infinitely large space with minimal reflections from the walls.

The methodology that the authors followed is actually very close in concept to that of Naples University [1, 2, 3] with the following differences: The experiments happen inside a fully Anechoic chamber and a Network Protocol Analyser (NPA) was utilised to generate the trigger for controlling the gate mode of the SA. In addition, we measure RSSI values from two wireless network monitoring devices. Therefore, in this work the Naples and the Maryland methods are combined to create a new measurement framework. This helps to conclude how accurate RSSI values are for measuring channel power at the PHY layer.

The following devices were used in the testbed: 1) A D-link DWL-900AP+ Access Point (AP). 2) A Toshiba Satellite Pro L300 laptop running Windows as Host 1. 3) A second Toshiba Satellite Pro L300 laptop acted as a wireless monitoring device running Backtrack Linux using an Asus Wi-Fi card with the Ralink rt73 chipset. This chipset was selected because it supports raw monitoring mode (rfmon), which allows wireless monitoring. The Wireshark tool was used for wireless frame capturing. 4) A Macbook laptop as Host 2. 5) A Netgear DS 108 hub. 6) A Fluke Optiview Series III Integrated Network Analyser with a wireless NIC. 7) A Spectrum Analyser (SA) Advantest R3182 connected to a 2.4 GHz dipole. 8) An Anritsu MD1230A Network Protocol Analyser (NPA).

Inside the full-anechoic chamber were located the AP, a laptop (Host 1), the monitoring devices: a laptop running Wireshark under Linux (with Asus Wi-Fi card) and the Fluke device, and a directional antenna as shown in Fig. 1. The Fluke provides Signal Strength and the Asus monitoring device provides RSSI values. The reason for using two devices was to compare different manufacturers tools because, as was discussed before, different manufacturers use different techniques to calculate the RSSI/Signal Strength information. In this case, we compare the Ralink rt73 chipset (Asus) with the Fluke Wi-Fi card.

A directional antenna (dipole) was attached to the SA in order to measure channel power. A dipole antenna radiates equally in all directions (isotropically) in a plane perpendicular to the dipole. However, the dipole has a null at either end. These can be thought of as blind spots where the antenna has difficulty receiving (or sending) signals. The dipole antenna in our experiments was vertically orientated to minimise the effect of the laptop (Host 1) on the measured channel power (see Fig. 2a).

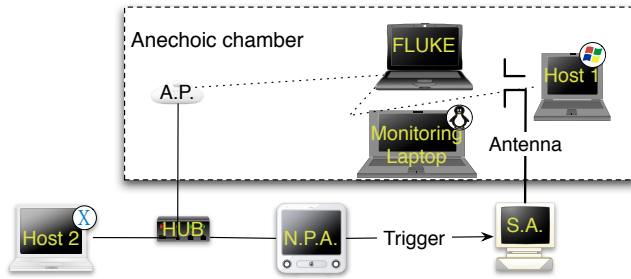In our experiments (see Section 4), the majority of the in-

Figure 1. A diagram of the testbed used.

formation was sent from Host 2 (via the AP) to Host 1. Host 1 sent back replies and acknowledgements. As Host 1 was much closer to the dipole than the AP, its contribution to the channel power would have been much larger if the dipole had been orientated horizontally or an omni-directional antenna had been used. A section of RAM was placed between the dipole and Host 1 (see Fig. 2 (a)). This reduced the channel power by 12 dBm in a specific test when only Host 1 was transmitting information.

Host 2 was connected with Ethernet to a hub, which propagated frames to the AP in the chamber and to the Anritsu NPA outside the chamber. The latter was set to monitor mode in order to monitor the frames received by the access point and generated triggers when several predefined conditions were met. In our case, the conditions were the specific IP source address (of Host 2) and destination IP address of Host 1. The trigger was used by the SA as a sign to start sweeping. Therefore, the SA only measured the channel power when a frame is sent from Host 2 to Host 1.
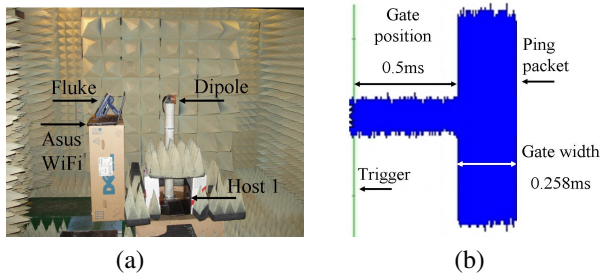

Figure 2. (a) A photo of the anechoic chamber with the equipment. The AP is behind the camera. (b) A photo of the oscilloscope displaying the trigger 0.5 ms before an icmp echo request packet.

3.2   Setting the gate sweep parameters

In wireless measurements of channel power, the frequency characteristics of frames are very important. However, usually the frame duration in a wireless network is very small, in fact smaller than the sweep time of the SA. This means that the SA will continue sweeping the desired frequency range (in our case Wi-Fi channel 6: 2.428 GHz - 2.448 GHz) after the frame has stopped being transmitted. The SA will not measure a signal inbetween frames and therefore the constructed frequency spectrum will be damaged as "holes" are introduced in the frequencies swept at this time (Fig. 3 (a)). This is a serious problem when measuring channel power. In order to circumvent this phenomenon, the spectrum analyser should be set in gate mode (Fig. 3 (b)). The gate mode is controlled by two parameters: i. Gate position: The time instance the SA should start sweeping - achieved with an external trigger. This is the delay between the NPA seeing the packet in the hub and the packet being actually modulated and transmitted from the AP. ii. Gate width: The duration of time for which the SA measures the signal. Gate mode allows one frequency sweep to be composed of several packets.
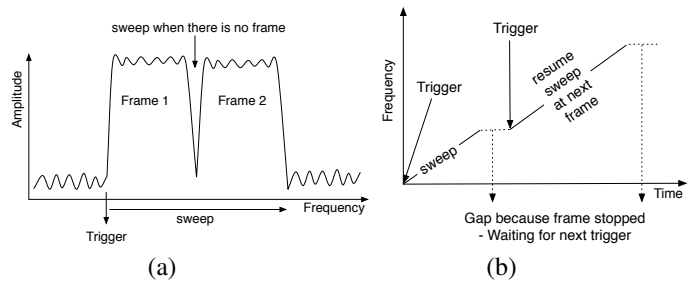
Figure 3. (a) The graph represents the display of a SA when the sweep captures instances when no signal is present (SA in non-gate mode). (b) In the gate mode, the SA receives a trigger and starts sweeping for a pre-defined time duration (gate width) and then pauses until another trigger is received, resuming the sweep of the subsequent portion of the frequency spectrum.

Before starting the practical experiments, the above gate parameters were examined for all three considered applications (ping, iPerf and iTunes). In order to achieve this, special alteration was made to the hardware of the AP. Specifically, an intermediate frequency (IF) signal at 380 MHz was taken from the AP to allow connection to an oscilloscope (DSO) for examination of the signal.

By triggering the DSO with the output from the NPA, we

239

Table 1. Parameters used in gate mode for each application

	Ping	iPerf	iTunes
Packet length in 802.11 (bytes)	120	1534	1536
Gate width (ms)	0.258	1.44	1.44
Gate position (ms)	0.5	1.06	1.06

could measure accurately both the gate position and gate width (see Fig. 2 (b)). The experiments showed that the trigger distance and the frame duration depend on the length (in bytes) of the packet that is modulated. Each of the three applications that are considered in this paper has constant packet lengths.

For ping; Host 2 sends echo requests. For iPerf; Host 2 is the client and sends UDP packets and for iTunes, Host 2 streams the video. The channel power from the spectrum analyser was measured using packets from Host 2 to Host 1 (as the trigger of the SA is generated under these conditions) and for averaging the RSSI/Signal Strength values of the monitoring devices, only packets from Host 2 to Host 1 were used.

For the ping application, we used the default ping size of 64 bytes (120 bytes as measured from the Fluke, i.e. 64 bytes + 20 for IP header + 8 for LLC + 28 for 802.11). For iPerf, UDP packets were 1470 bytes (1534 bytes in total as measured from the Fluke) and for iTunes the packets were 1472 bytes (1536 in total as seen from the Fluke). The parameters for each application are presented in Table 1. Note. without the gate mode, the channel power was reduced by 11 dBm, 7dBm and 1.5 dBm for ping, iPerf and iTunes respectively.

4 Practical Results and Discussion

A difficulty in performing these measurements is that ideally the receiving host (Host 1 in our case), dipole and monitoring devices should be in exactly the same location to ensure that the same power is measured in each case. Unfortunately, in practice this is impossible. Furthermore, if the objects are placed close together, they will interfere with each other electromagnetically. This is due to the dipole measuring unwanted signals emitted from the receiving host and to the reflections caused by having passive metal objects in close proximity to the antennas. Note, Host 1, Fluke and the monitoring laptop all contain antennas and their ability to receive signals may be reduced by surrounding objects, their orientation and due to the behaviour of the specific antennas. In the course of our experiments, we discovered

that if the RAM was placed in certain geometries it could adversely affect the performance of the antennas.

In this section, the cross-layer measurement results for ping, iPerf and iTunes are presented. The ping program on Host 2 generated 600 ping requests. iPerf was used as a UDP packet generator with a bit rate of 1 Mbps and a duration of 60 seconds. During this period of time, 5104 UDP packets were generated from Host 2. iTunes streamed two video clips to Host 1. Video 1 had a duration of 4:54 and Video 2 was 6:15 minutes long. Fluke and Asus averaged Signal Strength/RSSI values were compared against the measured channel power from the SA in order to examine their accuracy. In addition Fluke and Asus performance in wireless monitoring was compared.

Table 2 shows the results for ping. As the distance of the AP from Host 1 decreased, the higher layer results (average Round Trip Time (RTT) and duplicate packets) were improved. The channel power (Cp) measured from the SA, and the Fluke averaged Signal Strength increased as the distance decreased, which clearly reveals stronger signal intensity. Fluke also captured less retransmissions on the wireless link and slightly less frames in total at the shorter distance. Total frames for the Fluke device include frames from both directions of communication and retransmissions.

The RSSI results from the Asus Wi-Fi card agree with the Fluke and SA in terms of signal strength. However, the captured frames are much less in comparison to the number of captured frames from the Fluke. One reason for this is due to a possible driver - software - hardware compatibility problem with the specific instrumentation that did not capture frames generated from Host 1 and destined to the AP, thus limiting the possible capture capability to one way communication (AP to Host 1). Even though the capturing process was limited to one way, Asus still dropped a lot of frames that it was expected to capture.

For the iPerf application (Table 3), the SA channel power, Fluke's average Signal Strength and jitter follow the same behaviour as with the ping program, i.e. enhanced performance when the distance decreased. The Asus Wi-Fi card does not show increasing performance in capturing packets when the AP is closer. Surprisingly its averaged RSSI value remains very close to the value at 4.5 m. This clearly shows the prohibiting performance of specific Wi-Fi cards when accurate physical layer measurements are required due to the RSSI limitations.

Tables 4, 5 show the results for streaming two different video files in iTunes. The metric named Triggers refers to the number of packets seen by the NPA. These Trigger packets meet the specified triggering conditions. The number of Triggers is ideally the same as the number of packets that both Fluke and Asus should capture, ignoring packets from Host 1 and retransmissions at the MAC layer of the

240

AP. If all triggering packets are captured then the total capture frames of the monitoring devices will be greater than the number of triggers because of the captured retransmissions.

For both videos the experiment was run twice for the distance of 4.5 m and once for the distance at 1.5 m. Many other results have been produced but due to space limitations, Tables 4, 5 only include representative results and highlight the interesting features.

For Video 1 (Table 4), both experiments at 4.5 m generated similar channel powers. However, Video 1 was streamed for a longer duration in the second experiment (b) (1:41) in comparison to the first experiment (a) (1:24). This deterioration in the performance of the video can not be identified just by measuring channel power at the physical layer. The retransmissions metric taken at the link layer indicates the drop in the application layer performance that happened in the second experiment and results in the prolonged transmission of the video.

A similar phenomenon to the above can be seen for Video 2 (Table 5). The first experiment at 4.5 m (a) took a longer time to stream and had a higher number of retransmissions despite the channel power being higher than in experiment (b). Both Fluke and Asus capture more frames when the AP was closer despite the reduced number of retransmissions. It is notable that the Asus Wi-Fi card captures less packets than Fluke.

Asus' RSSI values do not indicate the real signal strength conditions as reported by the SA and the Fluke. Both Fluke and the SA indicated stronger signal strength at shorter distances as expected. However, the Asus RSSI decreases at shorter distances for both videos. The results for Video 1, Video 2 and iperf show that the RSSI values are neither reasonable nor intuitive. These practical measurements confirm the hypothesis that RSSI values are not reliable [5].

For Video 1 (Table 4), experiment (c) at distance 1.5 m took a longer time to stream in comparison to experiment (a) at 4.5 m. The number of retransmissions is not the only factor affecting the streaming time. The value of Triggers for experiment (c) indicates that more packets have been transmitted by Host 2 in comparison to experiment (a). This implies that the transport layer communication between the end hosts (Host 1 and Host 2) required retransmissions for several packets. Thus, the retransmissions from the AP to Host 1 (link layer) may be less than in case (a) but the retransmissions from Host 2 (transport layer) are greater. That causes the increased streaming time of Video 1. A similar phenomenon happens for Video 2 in experiments (b) and (c).

5 Conclusions and Future Work

Previously, a lot of research has focused on enhancing the already established protocols of network communication and to optimise them for a wireless channel medium. However, the underlying issues of the wireless medium and how these influence higher layers is not yet fully understood in the research community.

This paper formulates a framework for wireless network measurements and utilises a cross-layer methodology in order to address the limitations of traditional methodologies as discussed in section 2. The results of this paper demonstrate the importance of cross-layer measurements as a technique to identify the critical layers and interactions that are the culprits in the deterioration in the application layer performance. In addition, this paper compares current methodologies in terms of channel power measurement and relates these physical layer metrics to higher layers.

With the proposed methodology, the channel power is only measured while packets from Host 2 are transmitted and thus the precision is improved as packets from Host 1 and background noise are not included in the measurements. Using gate mode increased channel power by 11 dBm, 7dBm and 1.5dBm for ping, iPerf and iTunes respectively. In addition, the holes in the SA's frequency spectrum are reduced.

As expected, the practical results show that higher layer statistics improve as the distance between the AP and the client decreases. Generally, the performance, in terms of captured frames, of the two tested monitoring devices (Fluke and Asus) increased as the distance decreased. In addition, the RSSI results from Asus did not seem to give reliable results as they did not behave as expected and did not agree with the SA's channel power measurements nor with Fluke's Signal Strength. As was discussed in Section 2, the RSSI values are of no particular precision and are not intended for accurate channel power measurements. In this respect, Fluke behaves better than Asus and its Signal Strength seems intuitevly correct and agrees with the trends from the SA results.

As for future work, the trigger pulses that control the SA should ideally be generated from the AP that transmits and receives the frames on the wireless channel rather than from an external monitoring device (NPA in our case). This would result in the SA capturing retransmissions from the AP. In addition, for applications that vary their data packet lengths, a possible solution is the use of a state of the art DSO. A DSO captures the examined signal in the time domain and allows further processing of the signal information off-line. In this way, a DSO could calculate the power spectrum of a signal with varying pulse duration. However, this method may introduce severe restrictions in the time of the experiment due to memory limitations.

Table 2. Ping results

Distance (m)	4.5	1.5
SA Cp (dBm)	-55.81	-45.32
Ping Statistics:		
Packet Losses	0%	0%
Avg. RTT (ms)	22.55	11.99
Duplicates	11	3
Fluke Statistics:		
Avg. Signal Strength %	50.6	60.6
Total Retransmissions	14	8
Total Frames Captured	1207	1189
Asus Statistics:		
Avg. RSSI	69	80.7
Total Retransmissions	1	3
Frames Captured-1 way	254	464

Table 3. iPerf results

Distance (m)	4.5	1.5
SA Cp (dBm)	-55.95	-44.78
iPerf Statistics:		
Packet Losses	102	103
Jitter (ms)	6	5.63
Throughput (Kbits/s)	980	980
Fluke Statistics:		
Avg. Signal Strength %	44.5	59.2
Total Retransmissions	123	41
Total Frames Captured	3633	4922
Asus Statistics:		
Avg. RSSI	79.8	79.7
Total Retransmissions	106	29
Frames Captured-1way	2815	2731

Table 4. iTunes results for Video 1

Distance (m)	4.5	4.5	1.5
Experiment #	(a)	(b)	(c)
SA Channel Power (dBm)	-55.42	-55.84	-47.87
Triggers	26413	26506	26533
Time to stream	01:24	01:41	01:29
Fluke Statistics:			
Avg. Signal Strength %	49.26	48.43	56.1
Total Retransmissions	3440	6069	2124
Total Frames Captured	25934	29612	40685
Asus Statistics:			
Avg. RSSI	79.35	79.57	68.12
Retransmissions	696	1800	301
Captured (1 way)	6547	7400	10398

Table 5. iTunes results for Video 2

Distance (m)	4.5	4.5	1.5
Experiment #	(a)	(b)	(c)
SA Channel Power (dBm)	-54.6	-55.89	-48.27
Triggers	26762	26615	26960
Time to stream	01:38	01:14	01:18
Fluke Statistics:			
Avg. Signal Strength %	51.5	48.96	55.7
Total Retransmissions	4532	3092	2423
Total Frames Captured	27694	25635	41787
Asus Statistics:			
Avg. RSSI	79.25	79.3	68.66
Retransmissions	1354	630	279
Captured (1 way)	7731	6896	10294

References

[1] L. Angrisani, A. Pescape, G. Ventre, and M. Vadursi. Performance measurement of ieee 802.11 b-based networks affected by narrowband interference through cross-layer measurements. *Communications, IET*, 2(1):82–91, 2008.

[2] L. Angrisani and M. Vadursi. Performance assessment of wireless communication networks through cross-layer measurements. In *Instrumentation and Measurement Technology Conference, 2006. IMTC 2006. Proceedings of the IEEE*, pages 201–206, 2006.

[3] L. Angrisani and M. Vadursi. Cross-layer measurements for a comprehensive characterization of wireless networks in the presence of interference. *Instrumentation and Measurement, IEEE Transactions on*, 56(4):1148–1156, 2007.

[4] A. Balachandran, G. Voelker, P. Bahl, and P. Rangan. Characterizing user behavior and network performance in a public wireless lan. *Proceedings of ACM SIGMETRICS*, June 2002.

[5] J. Bardwell. You believe you understand what you think I said. . . , The truth about 802.11 signal and noise metrics. Technical report, Connect802 Corporation, 2007.

[6] C. Na. *IEEE 802.11 Wireless LAN Traffic Analysis: A Cross-layer Approach*. PhD thesis, University of Texas, May 2005.

[7] J. Yeo, M. Youssef, and A. Agrawala. A framework for wireless lan monitoring and its applications. In *Proceedings of the 3rd ACM workshop on Wireless security*, pages 70–79. ACM New York, NY, USA, 2004.

[8] J. Yeo, M. Youssef, T. Henderson, and A. Agrawala. An accurate technique for measuring the wireless side of wireless networks. In *International Conference On Mobile Systems, Applications And Services*, pages 13–18. USENIX Association Berkeley, CA, USA, 2005.