Title
Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis

Permalink
https://escholarship.org/uc/item/9jd3d55t

Authors
Cheng, Yu-Chung
Bellardo, John
Benko, Peter
et al.

Publication Date
2006-02-21

Peer reviewed

Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis

Yu-Chung Cheng, John Bellardo, Péter Benkö,*
Alex C. Snoeren, Geoffrey M. Voelker and Stefan Savage

Department of Computer Science and Engineering
University of California, San Diego

Abstract

The combination of unlicensed spectrum, cheap wireless interfaces and the inherent convenience of untethered computing has made 802.11-based networks ubiquitous in the enterprise. Modern universities, corporate campuses and government offices routinely deploy scores of access points to blanket their sites with wireless Internet access. However, while the fine-grained behavior of the 802.11 protocol itself has been well studied, our understanding of how large 802.11 networks behave in their full empirical complexity is surprisingly limited. In this paper, we present a system called Jigsaw that uses multiple monitors to provide a single unified view of all physical, link, network and transport-layer activity on an 802.11 network. To drive this analysis, we have deployed an infrastructure of over 150 radio monitors that simultaneously capture all 802.11b and 802.11g activity in a large university building (1M+ cubic feet). We describe the challenges posed by both the scale and ambiguity inherent in such an architecture, and explain the algorithms and inference techniques we developed to address them. Finally, using a 24-hour distributed trace containing more than 1.5 billion events, we use Jigsaw's global cross-layer viewpoint to isolate performance artifacts, both explicit, such as management inefficiencies, and implicit, such as co-channel interference. We believe this is the first analysis combining this scale and level of detail for a real production network.

1 Introduction

In the last five years, wireless networks based on the 802.11 family of standards have become ubiquitous in the enterprise. Integral wireless interfaces — now shipping in almost 90 percent of notebook computers — combined with unlicensed spectrum and inexpensive "access points" have made untethered Internet access a reality in almost two-thirds of U.S. corporations, hospitals and college campuses [11, 10, 6, 5]. However, the reality of these deployments can be quite complex. A large office building may have hundreds of wireless users interacting with dozens of access points under varying load and environmental conditions.

It is these interactions that make the dynamics of wireless network behavior so interesting, and yet so difficult to mea-

sure. Unlike their wired brethren, wireless networks are not well described as either a single broadcast channel nor as a graph of links. Whether any transmission is heard by a particular receiver is a function of many distinct factors including the ambient environmental interference, the sender's transmit power, the distance to the receiver, and the strength of any simultaneous transmissions on nearby channels as heard by that same receiver. Further complicating this morass is the 802.11 Media Access Control (MAC) protocol, which uses its own inferences about the physical layer to defer, schedule and retry transmissions. Finally, these networks are typically used to carry Internet traffic based on the TCP protocol that carries its own set of complex dynamics. It is no wonder that our understanding of these systems tends to be limited to either small controlled environments ("how much does interference between two radios impact throughput") or to large, but coarse measurements ("how long is the average TCP flow on a wireless network").

It is our belief that developing a deeper understanding of the dynamics and interactions in production wireless networks requires reconstructing their behavior in its entirety — measuring all frames and delivery outcomes across all wireless nodes. In the wired network this kind of measurement is typically achieved through passive monitoring, but in the wireless domain spatial diversity prevents any single monitor from capturing more than a small subset of traffic. Thus, extracting a global viewpoint requires many spatially dispersed monitors working in concert.

In this paper, we approach this problem from a systems point of view. We have developed a large-scale monitor infrastructure that overlays a building-scale production 802.11b/g network with over 150 passive radio monitors. These monitors in turn feed a centralized system, called Jigsaw, that uses this data to produce a precisely synchronized global picture of all physical, link-layer, network-layer and transport-layer activity. We believe our principal contributions are threefold:

- *Large-scale Synchronization.* We have designed and implemented a passive synchronization algorithm that can accurately synchronize over 150 simultaneous traces down to microsecond granularity. To accomplish this at scale requires predicting the impacts of individual radio clock skews and using every available frame

---

*Benkö is a visiting researcher at UCSD from the Traffic Analysis and Network Performance Laboratory (TrafficLab) at Ericsson Research, Budapest, Hungary.

1

to re-synchronize.

- *Frame Unification.* We use this fine-grained synchronization to combine the contents of all traces, merging duplicates and constructing a synchronized single trace of all frame transmissions.

- *Multi-layer Reconstruction.* From raw frame data we reconstruct a complete description of all link and transport-layer conversations. To address the problem of missing data we have developed a set of inference techniques to deduce the presence, time placement and even contents of missing data. Our analysis uses transport-layer information to resolve questions, such as frame delivery, that can be inherently ambiguous at the link-layer alone.

Since Jigsaw constructs a complete description at each layer, it is possible to relate actions at different layers that would otherwise be impossible. For example, in analyzing a client's transport activity, our TCP analysis will automatically identify a timeout and retransmission. However, by looking in Jigsaw's global trace we may find that the client did send an acknowledgment, but that it overlapped a period of broadband interference from a microwave oven that affected the client's AP (likely preventing the ACK from being bridged). While we have not yet completely automated this particular analysis, it is well within the capabilities of our system.

The remainder of this paper is organized as follows: In Section 2 we review the important aspects of the 802.11 MAC protocol and the related work in wireless LAN measurement. In Section 3, we describe our measurement infrastructure followed by a description of how traces are merged and synchronized in Section 4. Section 5 explains how link-layer and transport-layer viewpoints are reconstructed from raw frame data. In Section 7 we demonstrate Jigsaw's capabilities through a set of measurements that exploit its unique ability to provide a global wireless network perspective. Finally, Section 9 summarizes our overall conclusions with constructing and using this wireless monitoring infrastructure.

2 Background and Related Work

In this section we offer a brief tutorial in the operation of the 802.11 protocol followed by a description of previous 802.11 measurement research.

The 802.11 media-access control (MAC) protocol is a CSMA/CA variant that uses "virtual carrier sense" to support an RTS/CTS capability and to protect multi-frame exchanges. When a node wishes to send it first validates that the channel is clear. If the channel stays idle for a set period of time (DIFS) it transmits. Otherwise, it selects a random backoff time from $0 \ldots N$, and tries again. 802.11 provides a link-layer retransmission facility. Thus when a uni-

cast packet is sent, the receiving station is required to respond immediately with an ACK packet. If an ACK is not received within a preset timeout then the node doubles N, calculates a new (likely longer) backoff time and schedules a retransmission (retransmissions are indicated with a special bit in each frame header). Each frame carries a "duration" field that indicates the number of microseconds needed to complete the transaction (*i.e.*, including any acknowledgments that need to be sent) and each node will defer transmission until this time has passed. Special RTS and CTS frames are optionally used to ensure that any "hidden terminal" nodes will hear the reservation request. Frames are addressed using 48bit IEEE MAC addresses, although some frames (such as ACK, CTS and RTS) only specify the transmitter or receiver. Frames from the same transmitter are distinguished using a 12-bit monotonically increasing sequence number carried in each DATA frame. Special management frames (BEACON and PROBE) are used to discover the presence and capabilities of access points, while others (ASSOCIATION and AUTHENTICATION) are used to specifically connect a client to an access point.

802.11 supports a wide range of physical-layer implementations – the most popular being 802.11b (CCK modulation with coded rates up to 11Mbps) and 802.11g (OFDM, coded up to 54Mbps). Each client is responsible for choosing the rate to transmit each frame and this choice is encoded in the PLCP header at a "slow" rate (1-2Mbps for 802.11b, 6Mbps for 802.11g). However, "legacy" 802.11b radios are unable to decode the OFDM encoding of an 802.11g frame and can incorrectly sense the medium as idle. To address this problem, 802.11g access points determine if they have any 802.11b stations as clients. If so they enable "802.11g protection mode" in which each 802.11g frame is preceded by a low-rate CCK-coded CTS frame (CTS-to-self) that reserves the channel for the time needed to complete the 802.11g transaction.

Over the last 15 years, a progression of wireless network measurement efforts has provided insight into the behavior, performance, and reliability of 802.11 and precursor wireless LAN technologies. Early efforts used active measurements to study WaveLAN networks, one of a handful of early wireless LAN products in the 900 MHz band [7, 8, 21]. Since these relatively new wireless LANs were expected to have more severe error characteristics than wired LANs, an initial primary concern was the effect of wireless LAN errors on higher-level protocols and application performance. Consequently, these early efforts focused on evaluating the effects of errors on application performance, and analyzing and modeling packet and bit errors as a function of various factors such as distance, obstacles, and co-channel interference.

The pace of subsequent wireless measurement efforts increased as 802.11 technology was introduced and matured, and widespread deployments became commonplace. The

goals of these efforts similarly expanded to a broader range of concerns, exploring a wide range of environments, at increasingly larger scales, and with more extensive analysis over time. Measurement efforts explored university campuses [12, 13, 18, 19, 24, 25, 27, 28], industrial factories [26], corporate networks [4], and conference and professional meetings [3, 15, 16, 22, 23]. These efforts correspondingly became more extensive over time, from weeks of traffic from 75 users and a dozen APs in a department network [25] to years of traffic from thousands of users and hundreds of APs across an entire university campus [12]. These efforts also analyzed a wider range of characteristics of user behavior and network performance, such as application workloads, user session durations and user mobility, network installation and maintenance issues, error characteristics, *etc*.

Until recently, however, measurements of production 802.11 networks have treated them as a black box. For ease of methodology, these latter efforts typically have traced traffic on the wired distribution network and polled SNMP management data from APs as a basis for analyzing wireless LANs. As a result, such 802.11 measurement efforts have extensively characterized *what* user behavior and network performance wireless LANs provide, but have provided little insight into *why* applications and users experience such behavior and performance.

Recently, some researchers have started addressing this question by extending wireless network measurement to capture and analyze link-level characteristics as well. Such efforts require a change in methodology, however, since the measurement platform must observe raw network events. One approach is to use the same devices for experimentation as for observation. This approach works well for small, controlled active measurements, such as understanding link-level characteristics of outdoor mesh networks [2], error characteristics of factory environments [26], or components of handoff latency [20].

Passively monitoring the link-layer events of a large deployed network, however, requires dedicated wireless monitors separate from the wireless devices generating traffic. Three recent efforts have used this approach. Yeo *et al.* were the first to explore the feasibility of using separate monitors for passive wireless network measurement using synthetic experiments on an isolated 802.11 network [27, 28]. They use beacon frames to merge traces of a single flow observed from three wireless monitors, and demonstrate the utility of merging observations to improve monitoring accuracy. Jardosh *et al.* analyze the link-level behavior of traffic from a large IETF meeting using three monitors capturing traffic on orthogonal channels [15, 16]. They characterize and correlate retransmissions, frame size, and rate adaptation with reliability. Finally, a study by Rodrig *et al.* shares a number of the goals of our work [23]. They use five distributed wireless monitors to capture network events in a large conference venue. Using trace data from one of their monitors, they

characterize the extent of 802.11 management overhead and retransmissions, and analyze the effectiveness of rate adaptation for the clients in their trace.

Our work substantially extends previous efforts in wireless network monitoring in terms of scale, methodology, and analysis. Whereas previous efforts have used a small handful of monitors [15, 23, 28], our measurement platform uses over 150 monitors distributed throughout four floors of a 150,000 square-foot building to achieve extensive spatial and channel coverage. Tracing at such scale, however, presents new methodological challenges, such as globally synchronizing events in time across subsets of monitors as well as across channels; previous efforts either focus on separate channels [15], do not merge traces among monitors [23], or merge only a small number of traces using globally observed events [28]. Such extensive monitoring also presents new opportunities for analysis, in particular the ability to observe a large wireless network from a global perspective. From such a perspective, for example, we can analyze the extent and impact of co-channel interference in the network.

3 Data Collection

Any data analysis is ultimately predicated on the quantity, quality and precision of data that can be collected. While we believe that our analysis techniques are mostly generic, many of our design decisions *have* been informed by the capabilities of our infrastructure as well as the unique problems presented by its scale. For example, our approach to clock synchronization was driven by the need to merge data from 156 simultaneous traces, spanning a wide spatial and frequency range. In a smaller-scale environment a far simpler approach would have sufficed. Thus, to better motivate our constraints and opportunities, we use this section to describe our monitoring environment and the hardware/software infrastructure we have built to produce the raw traces for our analysis.

3.1 Environment

All of our measurement work takes place within the UCSD Computer Science and Engineering building — a large four-story structure shown in Figure 1. The building houses over 500 faculty, researchers, students and staff members within roughly 150,000 square feet with a total interior volume well over 1 million cubic feet. Production wireless service is provided by Avaya AP-8 access points (shown as triangles), which are configured to provide both 802.11b and 802.11g service[1]

Between and among these production APs we have deployed a constellation of 39 wireless sensor pods (shown

---

[1]In addition to the 39 production access points shown, the half-wing basement (not shown) houses five additional APs. We also occasionally observed signals from 46 additional authorized access points from nearby buildings and 22 rogue access points (mostly outside the building, but several inside).
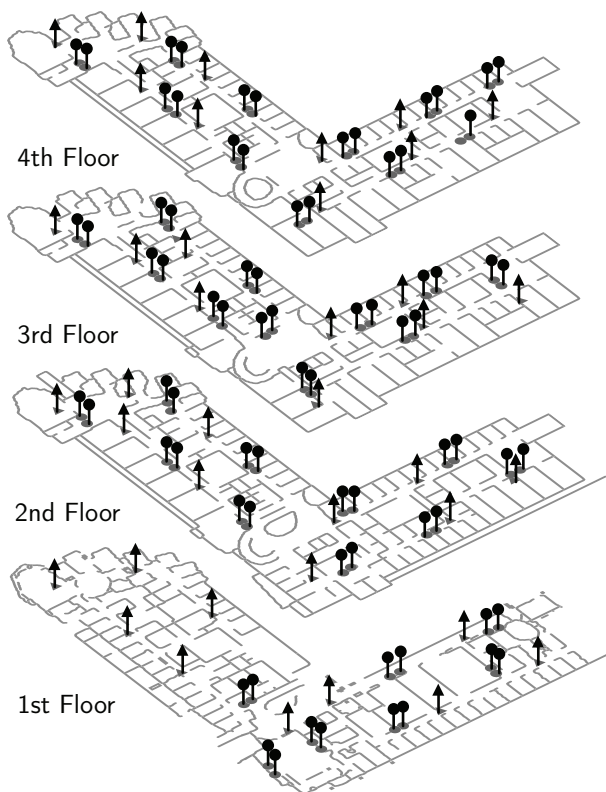
Figure 1: Building floorplan. This building comprises roughly 150,000 square feet spread over four floors (and a smaller basement, not shown). Circles indicate wireless sensor pods, and triangles indicate production access points.

as pairs of circles).[2] Each pod in turn comprises four independent radios, allowing for simultaneous monitoring at four distinct center frequencies – including all "non-overlapping" channels (1, 6 and 11) typically used in 802.11b/g deployments. The density of deployment, combined with this multi-channel capability, provides a "best case" scenario for capturing global behavior. We are unaware of any production wireless network monitored at similar scale.

3.2 Hardware

Concretely, each sensor pod consists of a pair of monitors set a meter apart. This organization provides sufficient antenna separation for active measurement experiments, while still being proximate enough to abstract both monitors as a single vantage point for passive monitoring. Each monitor is engineered from a modified Soekris Engineering net4826 system board, and pairs a 266-Mhz AMD Geode CPU with 128 MB of DRAM, 64 MB of flash RAM, a 100-Mbps Ethernet interface, and two Wistron CM9 miniPCI 802.11a/b/g interfaces based on the Atheros 5004 chipset. Each wireless

interface is connected, via shielded cable, to a separate external omni-directional "rubber duck" antenna mounted six inches apart on an aluminum enclosure. The antennas provide a signal gain of 2–3 dBi at 2.4 Ghz. Each monitor receives wired connectivity and power through a port on an HP 2626-PWR switch (seven in total).[3]

Finally, trace data from all 156 radios is sent via NFS to a single 2.8-Ghz Pentium server hosting 2 GB of memory and 2 TB of storage (four 500-MB SATA disks in a RAID-0 configuration).

3.3 Software

Each monitor runs a modified version of the Pebble Linux distribution designed for small memory embedded computers and a version of the open-source *madwifi* driver to drive the Atheros-based wireless interfaces [1]. We have made significant modifications to the driver to support additional transparency to the physical layer and improved capture efficiency.

Driver Modifications

We make three critical modifications to the default madwifi driver. While the standard driver only delivers valid 802.11 frames (even in so-called "monitor mode"), our version captures all available physical layer events, including corrupted frames and physical errors. Atheros hardware uses a 1μ resolution clock to timestamp each packet as it is received. Our driver slaves this timestamp facility to the clock of a single radio, so frames are recorded at both radios using the same time reference. Finally, our driver batches the delivery of physical event records, 64KB at a time, to amortize the impact of network load.

Jigdump

Data capture is managed by a specialized user-level application called *jigdump*. Each monitor executes two jigdump processes, one per radio, that are responsible for putting the wireless interface into monitor mode, "pulling" physical event records from the kernel and then transferring this data via NFS to a central repository. Jigdump reads data records 64 KB at a time via a standard `PF_PACKET` socket and generates an associated meta-data record that holds aggregate statistics and index information used to support subsequent random accesses. The raw data then is compressed using the LZO algorithm to minimize storage and I/O overhead (the two bottlenecks on our monitor platform). Data and meta-data are written to separate files via NFS, creating a new file pair each hour. In steady state, the NFS traffic across all 156 simultaneous feeds averages between 2–4 MB per second.

Together, jigdump and the driver modifications comprise roughly 950 lines of code.

---

[2]Our monitoring infrastructure does not cover the basement nor the northern wing of the 1st floor, which is not under our administrative control.

[3]Soekris Engineering uses an incompatible implementation of the 802.3af Power-Over-Ethernet standard and thus each system board is modified by hand to allow the HP switch to drive it.
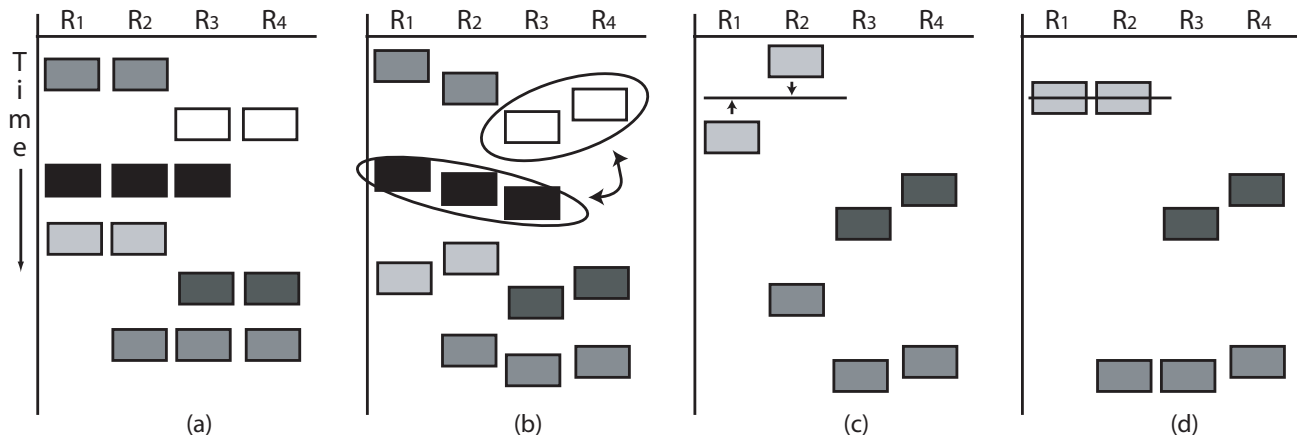
Figure 2: Description of synchronization steps, each color shade represents a unique frame. (a) "Wall clock" time at which frames received at four different radios, $R_1, ..., R_4$. (b) "Local time" at which frames received; bootstrap process groups identical frames until there are enough overlapping sets to normalize the time of every radio. (c) Later, clock skew causes timestamps for simultaneous frames to diverge. Timestamps for R_1 and R_2 are adjusted. (d) R_1 and R_2 are resynchronized and this process is repeated. Note that the adjustment of R_2 has also closely aligned subsequent frames with in-sync clock R_4 as well.

4 Trace Merging

Each individual trace represents a particular local vantage point on wireless activity. However, to construct a global viewpoint it is necessary to combine these traces into a single coherent description. This merging procedure must satisfy three key requirements:

1. *Unification*. A particular frame may be heard by multiple radios and therefore appear in multiple traces. It is important that these "duplicates" are identified as belonging to the same frame. In some cases a frame may not even be a perfect duplicate (*e.g.*, due to corruption or truncation), yet it is important that this not be treated as a unique frame.

2. *Synchronization*. While each frame is timestamped in each trace, the clocks used can vary significantly. To place these frames in proper order it is necessary to synchronize all frames to a common time standard. However, merely capturing the logical order is not sufficient for performing fine-grained analyses — like inferring interference between simultaneous transmissions. Such studies require all frames to be synchronized to at least the precision of a physical layer "slot time" (20 μs for 802.11b and 802.11g).

3. *Efficiency*. To permit on-line applications, trace merging should execute faster than real-time and scale well as a function of the number of traces. Thus, we desire an algorithm that can merge traces in a single pass over the data.

Our approach, similar to Yeo *et al.*'s framework [27], exploits the broadcast nature of wireless. Since wireless is fundamentally a broadcast medium, each transmission can potentially reach multiple in-range receivers. Moreover, in an indoor environment, propagation delay is effectively instantaneous — less than 1 microsecond to cover 500 meters at 2.4 Ghz. Consequently, we can treat the time at which a given frame is received by multiple monitors as a simultaneous instant for all potential interactions. Thus, we can use frames heard by multiple monitors as a common reference point to synchronize the clocks at each monitor and globally order subsequent events between traces. Finally, these reference frames can be used to calculate global timestamps for subsequent events *within* each trace, using the offset in the local clock to place them accurately. Subsequently, identical frames with the same timestamps can be unified, thereby creating a single global trace. In the remainder of this section we describe the synchronization and unification algorithms used by Jigsaw.

Our synchronization approach is inspired by Elson *et al.*'s RBS protocol for sensor networks, which shares many of the same assumptions [9]. The two algorithms, however, diverge significantly in implementation due to the differing demands of their applications: Jigsaw must be opportunistic in finding time references yet permits a centralized implementation, while RBS mandates reference broadcasts, yet requires a distributed implementation. Most importantly, RBS provides relative time synchronization between pairs of sensors, while Jigsaw must accurately synchronize all traces to a single global clock. Accomplishing this involves two phases: bootstrapping the synchronization algorithm to instantiate a single universal time standard across all radios, and then maintaining synchronization to this standard during frame unification.

4.1 Bootstrap Synchronization

Bootstrapping is accomplished by finding reference points to synchronize the radios of a set of individual monitors

and then synchronizing between sets until a single — albeit imaginary — coordinated time standard is established.

More precisely, let r_i denote the ith radio and let T_i represent the difference between its clock and "universal time" — the global time reference we hope to agree on. Let s_k denote the kth reference frame used to synchronize radios and let $E_{i,k}$ be the set of pairs $<r_i, s_k>$ such that radio r_i receives frame s_k. Moreover, let y_{ik} denote the local value of r_i's clock when it received s_k (defined if and only if $<r_i, s_k>$ is in E) and T_i to be the offset needed to adjust r_i's clock to universal time. Thus, when s_k has been received, the universal time can be defined as

$$U_k = y_{ik} + T_i.$$

To bootstrap synchronization, Jigsaw must simply find T_i for each radio. Once the offset T_i is available, Jigsaw can place each frame s_ks into universal time by adjusting its timestamp y_{ik}.

Ideally Jigsaw could locate a *single* 802.11 frame s_k received by all radios and then y_{i1} could be picked arbitrarily to represent the initial universal time. Unfortunately, we cannot depend on such events in a large deployment since signal strength decays with distance and no single frame is likely to cover an entire building. Moreover, real deployments use multiple channels and a frame transmitted on one channel may never be heard by a monitor on another.

To overcome this problem, we can synchronize transitively via overlapping subsets of radios that are each synchronized with each other. For example, suppose radio r_1 and r_3 are too far apart to share any reference frames, but each share distinct reference frames with an intermediate radio r_2. If s_1 is a reference frame received only by r_1 and r_2, and s_2 is a reference frame only received by r_2 and r_3, then $y_{1,2}+T_1 = U_1 = y_{2,1}+T_2$, and $y_{2,2}+T_2 = U_2 = y_{3,2}+T_3$. Then $T_3 = y_{1,1} - y_{2,1} + y_{2,2} - y_{3,2} + T_1$. The more densely radios are deployed, the more such transitive paths between r_1 and r_3 are likely to exist. However, to maximize the likelihood that T_is are globally consistent — meaning that $(T_j - T_i)$ plus $(T_k - T_j)$ will equal $(T_k - T_i)$ — we try to maximize the overlap between paths by minimizing the number of distinct reference frames.

Our protocol works as follows. Jigsaw examines the first second of data from each trace.[4] For each frame s_k, Jigsaw checks if it was also received by any other radios. If Jigsaw finds an identical frame heard by some radio r_i, it adds r_i into E_k. Note that not all 802.11 frames are good references for synchronization. For example, ACK frames to the same destination are always identical, the same stations are always identical, some stations always use zero sequence numbers

on probe frames, and frame retransmissions cannot be distinguished from one another. Thus, Jigsaw only uses "unique" frames for all synchronization activities. Generally, these are DATA frames that do not have the retransmit bit set.[5]

For every radio trace, Jigsaw picks the set E_k that contains the maximum number of radios and adds it into the synchronization set G. Jigsaw stops filling G when G contains an instance of each radio. At this point, there usually exists at least one path between any arbitrary two radios (if not, the original one-second window could be widened or more radios added to G, but we have never had need to do this). Then, for each radio r_i, Jigsaw performs a depth first search in G to reach r_1 (see parts (a) and (b) in Figure 2). Recently, Karp *et al.* [17] have discussed ways of picking the optimal paths for a similar problem, but we have found that most paths from r_1 to r_i are precise enough in practice ($\pm$ 10 μs).

As described, this algorithm is sufficient to synchronize all radios on the same channel. However, there is no transitive path between radios on strongly disjoint channels. To fully synchronize the trace we exploit the fact that our monitors use a single clock to timestamp frames received on both of their radios. Thus, in this particular context local timestamps for frames on one channel can be directly related to timestamps on another — effectively bridging a path between them. However, since this particular set only contains two radios, by definition it is unlikely to be picked early in constructing G (leading to excessively long paths and poorer accuracy). Therefore, we modify our algorithm to prioritize the use of inter-channel sets early in the construction of G.

4.2 Frame Unification

After bootstrap synchronization, Jigsaw processes all traces in time order and unifies duplicate frames, called *instances*, into a single data structure called a *jframe*. Each jframe holds a timestamp, the full contents of the frame and the identity of the radios that heard each instance. Figure 3 provides an example of this source data as it is being unified. As part of the unification process, Jigsaw also aggressively resynchronizes the clocks between each trace. We describe the evolution of our algorithm below.

Basic Unification

For each trace (*i.e.*, radio) Jigsaw maintains a frame queue sorted in time order. The simplest unification approach is to linearly scan the head of all frame queues and group the frames with the same timestamps and contents. More concretely, Jigsaw will select the first valid frame (*i.e.*, FCS was successful) as the representative instance and then perform content comparisons to find instances among the candidates. To quickly prune false negatives, Jigsaw compares frame length and FCS fields first and short-circuits the comparison on failure. For PHY-restart frames or frames with invalid

---

[4]In this case, "the first second" refers to true time (UTC) as measured by the system clock on each monitor. Each monitor maintains their system clock within milliseconds using the NTP protocol and records this value in its traces. This is the only point at which the system clock time is ever used.

[5]Some Intel 802.11 implementations incorrectly retransmit data *without* the retransmit bit set, but thankfully this is rare.

Jigsaw

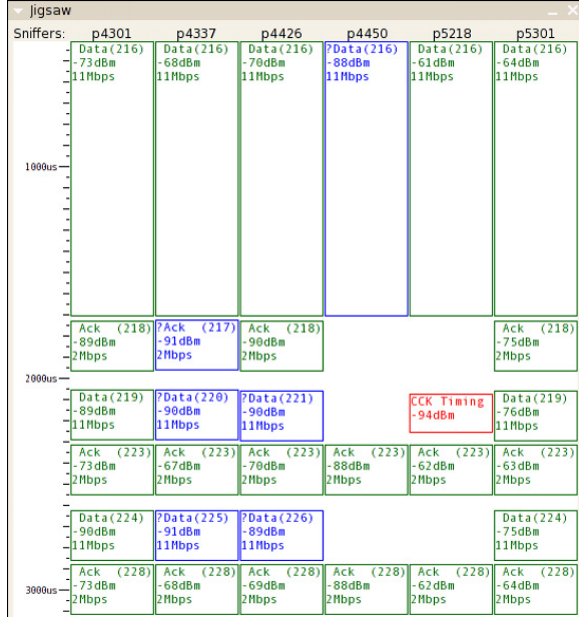| Sniffers:	p4301	p4337	p4426	p4450	p5218	p5301
	Data(216) -73dBm 11Mbps	Data(216) -68dBm 11Mbps	Data(216) -70dBm 11Mbps	?Data(216) -88dBm 11Mbps	Data(216) -61dBm 11Mbps	Data(216) -64dBm 11Mbps
	Ack (218) -89dBm 2Mbps	?Ack (217) -91dBm 2Mbps	Ack (218) -90dBm 2Mbps			Ack (218) -75dBm 2Mbps
	Data(219) -89dBm 11Mbps	?Data(220) -90dBm 11Mbps	?Data(221) -90dBm 11Mbps		CCK Timing -94dBm	Data(219) -76dBm 11Mbps
	Ack (223) -73dBm 2Mbps	Ack (223) -67dBm 2Mbps	Ack (223) -70dBm 2Mbps	Ack (223) -88dBm 2Mbps	Ack (223) -62dBm 2Mbps	Ack (223) -63dBm 2Mbps
	Data(224) -90dBm 11Mbps	?Data(225) -91dBm 11Mbps	?Data(226) -89dBm 11Mbps			Data(224) -75dBm 11Mbps
	Ack (228) -73dBm 2Mbps	Ack (228) -68dBm 2Mbps	Ack (228) -69dBm 2Mbps	Ack (228) -88dBm 2Mbps	Ack (228) -62dBm 2Mbps	Ack (228) -64dBm 2Mbps

Figure 3: Jigsaw visualization of synchronized trace. Time appears on the x-axis in microseconds and individual radios (only six shown here) on the y-axis. At roughly 400μ a client sends a data frame that is heard by all six radios. However, radio "p4450" is too far away (signal strength of -88dbM) and both the frame is corrupted and the subsequent ACK is not received. However, more than enough radios are present to construct a jframe for path parts of the frame exchange. At 2000μ a different client sends and it is heard by a different set of radios. Note that p5218 is too far away to even synchronize with the preamble.

FCS, Jigsaw cannot perform a full content comparison and simply matches on the transmitter's address field (since these frames are not directly used for any higher-layer reconstruction, any rare false matches will have little impact).

However, there are two problems with this approach. First, for large deployments the linear scan can have tremendous overhead. In our environment, most jframes contain 10 or fewer instances and yet we have over 150 simultaneous traces whose queues must be checked. To minimize this overhead, Jigsaw instead populates a single priority queue sorted by time with the most recent frame from each trace. To create a jframe, Jigsaw simply pops this queue until the timestamp changes and groups the resulting candidate frames according to their content (it is still crucial to compare frame contents since it is possible that distinct frames may be serendipitously transmitted at the same time). Thus, the time to create a new jframe is roughly $O(\log(N))$ instead of linear time.

The second problem is that each radio's clock skews over time. The 802.11 standard mandates an accuracy of at least 100 PPM (0.01%) and our experience is that Atheros hardware has far better frequency stability in practice. However, even good clocks diverge from each other over time. If the time offset between clocks becomes great enough, then some instances of a given frame may not be correctly merged into the same jframe. To mitigate this problem, we modify the procedure for creating a jframe. After the first frame candidates are popped and grouped, we remove additional frames from the priority queue until the timestamp at the head of the queue exceeds some *time offset threshold* with respect to the candidate instances — *i.e.*, a "search window." Some of these additional frames will have identical content with the other candidates and will be grouped into a new jframe while the others will simply be reinserted back into the priority queue.

Clock Adjustment

Solving this unification problem also provides a means to resynchronize the traces. When a set of frame instances are unified the time differences between their timestamps represent how much each clock now differs (again, it is critical that we only use unique frames to drive this synchronization). We select the median timestamp value and elect it to become the new universal time reference. The difference between this value and the timestamps on each instance represents a correction factor — positive or negative — that is then used to bring each of the associated traces back into synchronization (see parts (c) and (d) in Figure 2). A tradeoff can be made between accuracy and the overhead of resynchronizing by placing a threshold on the minimum *group dispersion* — the difference between the earliest and latest timestamp for a frame instance — before resynchronizing. In our implementation we set this threshold to 10 μs. (Note that this does not limit the synchronization accuracy to 10 μs.)

Managing Skew and Drift

If resynchronization happened frequently and uniformly across all traces, then it would be easy to maintain very tight synchronization bounds. However, there are frequently extended periods (although rarely over 100 ms since this is roughly the period between AP beacon frames) during which a particular radio may not observe any frames in common with others. During these times the synchronization of this radio's observations is only guaranteed by the accuracy of its own local clock. Thus, the slope of its skew with respect to universal time will determine how quickly it will lose synchronization without readjustment. In practice, we have found that with large numbers of radios, unless the search window is made dangerously large (100s of milliseconds) perfect synchronization is lost quickly. However, many of these problems can be eliminated by incorporating measurements of per-radio clock skew into the synchronization algorithm. Thus, the timestamp of each frame is adjusted to compensate for the impact of the clock skew on the radio receiving it. In addition, for large numbers of radios we have also found it important to compensate for clock *drift* — the change in skew over time — by using an exponentially weighted moving average of past skew measurements to predict future skew on a per-frame basis.

Figure 4: Cumulative Distribution Function of group dispersion across all frames.

Thus, Jigsaw can use almost every new data frame for continual resynchronization.[6] This presents several key advantages compared to approaches that simply use reference beacons to synchronize [27]. First, in large environments it is not possible to identify frames heard by *all* monitors and thus time synchronization must be transitive. Having more synchronization actions will almost always increase synchronization accuracy since the impact of clock skew is minimized. Second, since clients are mobile, their traffic creates a richer set of synchronization opportunities — touching pairs of radios that might never be directly synchronized otherwise. Finally, more clock samples allow for better management of skew and drift and therefore accuracy. In small-scale environments these factors may be minor. As the number of monitored radios increases, however, variability in skew, drift and workload conspire to raise the probability of a synchronization loss. This additional robustness becomes critical at a modest increase in complexity. Jigsaw's synchronization and unification code totals roughly 4,000 lines of C++.

Figure 4 illustrates the current accuracy of our algorithm using a 10ms search window over unique frames. The graph shows the CDF of group dispersion values calculated for every frame processed from 156 radios over a 24-hour period. Thus for 90% percent of all frames, the worse case time offset between any two radios is less than 10 μs, and 99% see a worst case offset under 20 μs. While the details of this graph are a function of individual clock characteristics, the network workload and the number of clocks being kept in sync, we believe it demonstrates that fine-grained broadcast synchronization is achievable in a building-scale environment.

5   Link and Transport Reconstruction

Having constructed a single global view of each observed physical event, the next task is to reconstruct each link-layer

---

[6]In truth, there are still a few esoteric reasons why synchronization can be briefly lost. However, we have encountered these problems in less than one thousandth of one percent of over 530M frames processed in our trace.

and transport-layer conversation in its entirety. In principle, this is easy since Jigsaw provides a time ordered list of all frames and each frame contains up to 200 bytes of payload that can be used to identify MAC addresses, IP addresses and TCP port numbers. However, in practice this construction is complicated by missing data and by vantage point ambiguities. Thus, Jigsaw must use inference to help reconstruct the these higher-layer descriptions.

5.1   Link-layer Inference

In reconstructing link-layer conversations, Jigsaw first identifies each *transmission attempt* from a sender. For example, a CTS-to-self packet, a subsequent DATA frame and the trailing ACK response may all be part of the same attempt. To group these together automatically we first use the MAC address: DATA frames carry the address of the sender explicitly, CTS-to-self frames (used for 802.11g protection) do as well and ACK frames indicate the recipient's address. Then we examine the *Duration* field in the CTS and/or DATA frames, and use it to deduce the future time in which an ACK, if sent, must have been received. This timing analysis is especially critical when frames are missing from the trace since otherwise we might risk assigning an ACK for a missing DATA frame to an earlier observed DATA frame.

We then group transmission attempts into *frame exchanges* — complete sets of transmission attempts that end in a link-layer frame being successfully delivered or not. Since 802.11 implements ARQ for unicast frames, a frame exchange may involve multiple transmission attempts. Normally it is sufficient to simply group nearby transmission attempts that share the same frame sequence number. However, when portions of transmission attempts are missing (*e.g.*, CTS and ACK, but not DATA), then we must deduce the presence of or absence of this missing data based on the subsequent behavior of the sender and receiver. For example, if we observe a lone CTS-to-self frame immediately followed by a subsequent CTS/DATA frame pair with the retry bit set, we can infer that the first CTS was followed by a DATA frame (and indeed that this frame held exactly the same content as the latter). Moreover, we can infer the rate at which this packet was sent based on the length of the subsequent packet and the size of the duration field in the first CTS. Finally, we *hypothesize* that the first data packet was lost since we did not observe an ACK and acknowledgments are lost less frequently than data. Our inferences are implemented using a finite-state machine capturing the visible aspects of the transmitter's MAC state in addition to several heuristics (*e.g.*, that DATA is more likely lost than ACKs). We do not make inferences about frames for which we have no direct information (*i.e.*, sequence gaps greater than one) but our experience is that these situations occur rarely in our traces.

Finally, one of the most important questions we wish to infer is whether a particular frame exchange was successful

or not. Unfortunately, the vantage point of a passive monitor does not allow this to be determined unambiguously. To wit: if, after transmitting a DATA frame, we see an ACK, we can feel confident that the data was delivered. However, if we never see an ACK, it is ambiguous if the frame was lost or if we simply did not observe the ACK. However, we *can* disambiguate this situation by using transport-layer information.

5.2 Transport Inference

Our transport-layer analysis takes frame exchanges as input and reconstructs individual TCP flows based on the network and transport headers. We use a variant of Jaiswal *et al*'s analysis (designed for wired passive monitors) to then infer connection characteristics (*e.g.*, RTT, RTO, fast retransmissions, segment losses, *etc.*) [14]. The passive wireless context, however, has two ambiguities that differ from the wired environment. First, we may process frame exchanges in which it is unclear if the frame was actually delivered or not (as described previously). However, we can frequently use the transport-layer side effects of this frame as an oracle to determine what truly happened. For example, a data frame carrying a new TCP segment will cause subsequent TCP acknowledgments to "cover" its TCP sequence space. Thus, observing a covering TCP ACK *proves* that the link-layer frame containing the associated data was actually delivered. To our knowledge we are the first to exploit transport-layer inference to ascertain link-layer delivery. The second problem is that existing analyses assume that monitors are lossless (that is, they observe all packets that are delivered between endpoints). In the wireless content, even with many different monitors, sometimes a frame exchange is completed but not observed *at all* by a monitor. Thus we have modified Jaiswal *et al.*'s analysis to infer the delivery of unobserved TCP data based on protocol behavior (our analysis is robust to any single loss).

6 Validation

A fundamental challenge with distributed wireless monitoring is obtaining effective coverage of all network transmissions. Since the monitors are not co-located with either clients or APs, it is possible for monitors to miss some network transmissions due to range, environmental conditions, interference, *etc.* In this section we describe two experiments to evaluate the coverage of our monitoring platform at both the link and transport layers.

First, we performed a controlled experiment to compare link-level events measured using an observer with perfect knowledge with measurements using the monitor platform. Using a wireless laptop, we generated a network workload at various locations throughout the building. The workload was a combination of Web browsing on the Internet, interactive `ssh` sessions to wired hosts, and `scp` copies of large files. This workload produced both short and long flows as well as small and large packets. We generated this workload
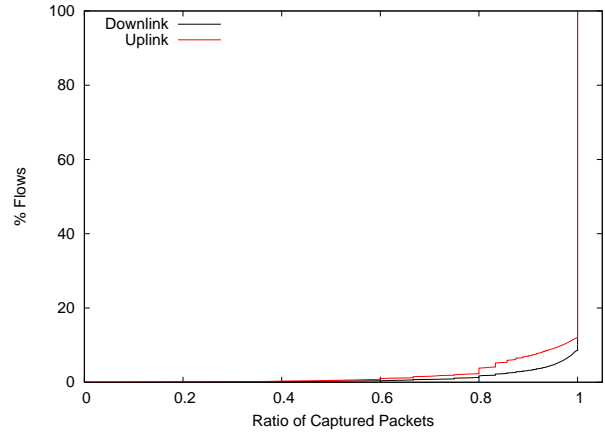

Figure 5: Coverage of packets in TCP flows by the monitoring platform.

at three locations in each wing of each floor. These locations ranged from heavy to light places of wireless usage.

While performing communication, the laptop recorded the link-level events it generated and observed from its associated APs. At the same time, we used the monitoring platform to observe the laptop communications. Comparing the events recorded at the laptop with those recorded by the monitoring platform, the platform observed 95% of all link-level events generated by the laptop. The coverage in this experiment is consistent with other studies using similar wireless monitoring methodology: [15] reports a coverage of 80–97%, [23] reports 90%, and [28] reports 97%.

Second, we compared the TCP flows captured in a day-long trace of the wireless network (described in more detail in the next section) with a second trace of the same traffic captured on the wired distribution network. We restricted the comparison to the set of TCP flows that could be possibly observed at both vantage points; the monitor for the wired network, for example, does not see traffic sent from one wireless host to another. For every packet in every TCP flow in the wired trace, we checked to see if the packet also appeared in the wireless trace. Overall, the coverage of TCP packets is excellent. For 57,782 TCP flows containing a total of 6 million packets in the wired trace, 98% of those packets also appear in the wireless trace. This high coverage is particularly encouraging since the trace includes distant clients connected to the building APs from the basement and the administrative wing on the first floor, locations lacking monitors.

Figure 5 shows the results of this experiment in more detail. Across all TCP flows, it shows the fraction of packets in a flow that appear in the wired trace that also appear in the wireless trace. It separates the flows into the downlink (to wireless network) and uplink (from wireless network) directions. The graph shows that, for most TCP flows, the monitoring platform captured all of their packets (88% uplink,

Start	1/24/06 @ 00:00
Duration	24 hours
Sniffers	156
Total APs	107
Our APs	39
Other APs	68
Our Clients	1,026
Total Events	2,700 M
Physical Errors	338 M (13%)
CRC Errors	956 M (35%)
Valid Frames	1,410 M (52%)
Jframes	530 M
Jframe Events	1,580 M
Events/Jframe	2.97

Table 1: Summary of trace characteristics.

91% downlink). Although the remaining flows have missing packets in the wireless trace, almost all still have high coverage. We also see a slight coverage difference based upon direction: the platform captures a slightly higher fraction of downlink packets than uplink packets. This effect is due to the fact that our monitors have better coverage of APs (*e.g.*, the monitors are closer to APs than clients on average).

Based upon the coverage measured in both experiments, we conclude that the monitoring platform provides sufficient coverage to perform detailed analyses of traces captured using the platform.

7 Analyses

In this section we perform a series of analyses on a trace of the building wireless network captured by the monitor platform. Since the amount of data and range of possible analyses is quite large, we focus on analyses that take advantage of the global perspective afforded by the distributed monitors. Our goal is not to be exhaustive, but rather to illustrate the unique capabilities of a global synchronized viewpoint and cross-layer analysis. We start by summarizing high-level characteristics of the trace, and then examine the effects of interference, the effects of 802.11g protection mode in networks with both 802.11b and 802.11g clients, and distinguishing link-layer and wired effects on TCP loss rate.

7.1 Trace Summary

We start by summarizing the high-level characteristics of our trace and then show network activity over time. Table 1 presents the characteristics of the trace we use for our analyses. The trace captures traffic for the entire day of Tuesday, January 24, 2006, a typical workday in our building. Just as APs within buildings are not isolated, buildings themselves are not isolated: we observe traffic associated with more than twice as many APs in surrounding buildings than in the building. For the subsequent analyses, though, we focus only on the traffic generated by clients associated with

our APs; our monitors cannot capture traffic from external APs with reasonable coverage due to their remote location. We see 1,026 unique client MAC addresses associated with our APs during the day.

Throughout the day the monitors observe over 2.7 billion events. Over 47% of these events are physical or CRC errors. This high percentage is not too surprising given transmissions observed by distant monitors just beyond reception range, the presence of both co-channel interference (hidden terminals) and broadband interference (microwave ovens), *etc.*

Jigsaw unifies 1.58 billion events (valid frames and a subset of associated error frames) into 530 million jframes, for an average of 2.97 events per jframe. In other words, on average the monitoring platform makes three observations of every observed transmission of a valid frame in the network.

Figure 6 shows network activity as a time series throughout the day at the granularity of one minute. Figure 6(a) shows the number of active clients and APs per one-minute time slot as a stacked bar graph. We define an active client as one that is communicating with an AP or is actively establishing an association. An active AP is one communicating with an active client (an AP only sending out beacons, for example, would not be active). Activity exhibits an expected diurnal pattern. Most clients are active from late morning (10am) until late afternoon (5pm), with many clients active in the early morning and well into the night. The number of active APs grows as more clients become active throughout the building. The clients active overnight are likely active wireless devices without user activity, such as wireless laptops left running with applications that produce background traffic.

Figure 6(b) shows the amount of traffic per one-minute time slot as a stacked bar graph of four traffic categories. "Data" counts both unicast and broadcast data frames, and "Management" counts various management and control traffic (RTS/CTS, ACKs, association, *etc.*). Although the number of active clients is relatively smooth over time. the traffic generated by those clients is much more bursty. Many of the bursts start on an hour or half-hour time boundary, likely indicating laptop usage during meetings and talks in the building. Since most management and control traffic relates to data traffic, it closely tracks the amount of data traffic.

We also separate out two explicit categories of management traffic because of their high prevalence: "Beacon" shows the amount of periodic AP beacon traffic, and "ARP" shows the amount of ARP broadcast ARP traffic. Because APs broadcast beacon traffic independent of activity, beacon traffic is constant throughout the day. ARP traffic is more interesting. In addition to legitimate use, outside scans and worms generate ARP traffic as they probe unallocated IP address space. However, it appears that the largest source of ARP is due to an 802.11 management server from Vernier that uses regular ARPs to track the liveness and network lo-
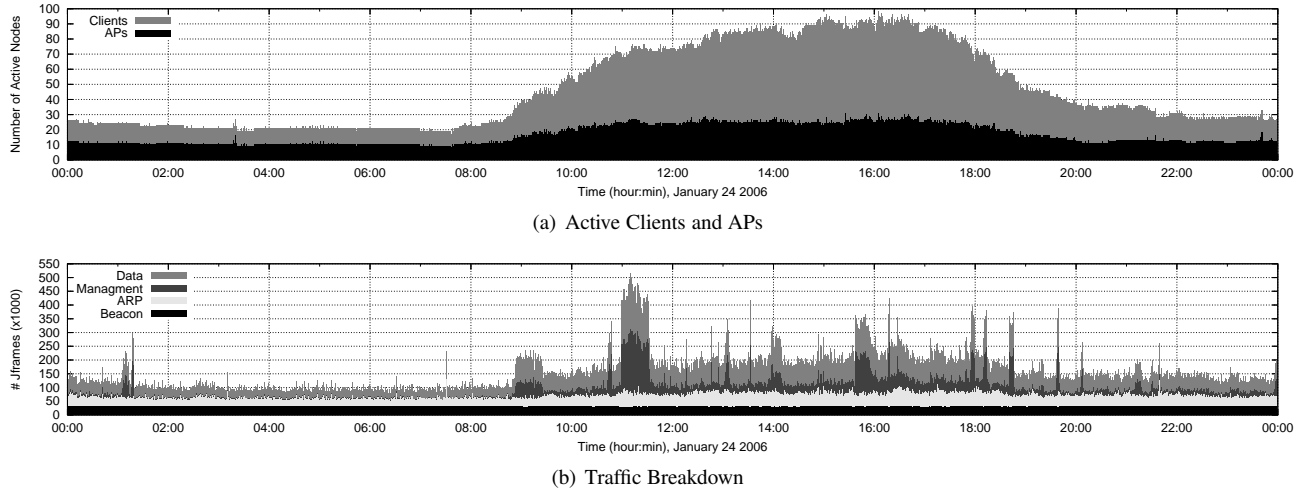
(a) Active Clients and APs

(b) Traffic Breakdown

Figure 6: Time series of network activity throughout the day in one-minute intervals.

cation of registered clients. However, the important aspect of ARP traffic is that it is broadcast. Because 802.11 APs are designed to act as transparent bridges all ARP "who-has" broadcasts from the wired network are also broadcast on the wireless channel. Since broadcast frames are always encoded at the lowest rate they make highly inefficient use of the medium. Indeed, if we examine our trace strictly from an *air time* perspective, broadcast traffic (primarily ARP and Beacons) regularly consumes 10% of the channel as seen by any given monitor. Finally, because they are delivered to all APs at the same time, they are broadcast on all APs on all channels at roughly the same time as well – likely interfering with themselves in the process.

Indeed, all network-layer broadcast traffic has this side effect as well, including client DHCP requests and application broadcasts [7] Moreover, aspects of this traffic scale with the size of the network or the size of the user population while the capacity of the channel remains constant. Thus, we argue that applications should use multicast instead of broadcast on 802.11 networks and 802.11 APs should be modified to perform selective filtering of non-unicast traffic. Finally, to eliminate the implicit synchronization caused by wired broadcasts, APs should add random jitter to the transmission time for broadcasts frames received from the wired network.

7.2 Interference

In this section, we analyze the extent of transmission interference experienced by nodes in our trace. Since the platform monitors orthogonal channels, adjacent-channel interference is rare and co-channel interference from hidden terminals is likely the dominate cause of interference. As a re-

sult, the distributed monitoring platform provides the key ability to observe co-channel interference. By providing a global perspective on the network, we can simultaneously detect a transmission from a sender to a receiver, hypothesize that the transmission was lost, and detect that a third node was transmitting at the same time as the sender. With only a single vantage point, it would be very difficult to detect and correlate such simultaneous transmissions.

We define an interference event as a unicast transmission from a sender s to a receiver r in which one (or more) interferers i simultaneously transmit and cause the transmission from s to r to fail. Based upon events in the trace, our goal is to estimate what fraction of these simultaneous transmissions cause a loss due to interference. Note that packet transmissions are distinct from frame exchanges; a successful frame exchange might experience multiple transmission losses and recover using link-level retransmissions.

We measure simultaneous transmissions when the trace contains more than one transmission overlapping in time during which s transmits a packet to r. We infer that the transmission from s failed to reach r when we do not observe an ack from r. At this point, though, when a loss happens we cannot say for certain that a particular simultaneous transmission was the true cause of the loss. It may be the case, for instance, that a node in a remote part of the building just happened to have transmitted at the same time as a transmission from s to r was lost; the loss itself may have been caused by any number of reasons entirely unrelated to the remote node's transmission.

We can, however, infer when losses are likely due to simultaneous transmissions. In particular, we can infer the conditional probability P_i of a simultaneous transmission causing interference given that there is a simultaneous transmission from s to r. We can infer P_i based upon the losses between s to r when simultaneous transmissions both do and

---

[7]One particularly egregious example – almost 100,000 frames in our trace – is the Mac version of the MS Office suite. As part of an anti-piracy mechanism the software regularly broadcasts its license information to UDP port 2222.

11

do not occur. Informally, if we assume that the background loss rate is constant regardless of the number of transmissions, we can attribute the losses between s and r during simultaneous transmissions accordingly: If s and r experience few losses in the absence of simultaneous transmission, the more likely the losses they experience during simultaneous transmission are due to interference.

More formally, let I be the event that interference causes a lost transmission from s to r, and L be the event that the transmission from s to r was a background loss due to some other cause (*e.g.*, range, obstacles). Let S be the event that there is a simultaneous transmission from at least one other device i when s transmits to r. Note that I and L are independent events. For the case where no multiple simultaneous transmissions occur, $P[I|\neg S]$ is obviously 0. Unfortunately, when there are multiple transmissions we cannot empirically distinguish between I, L, or $(I \cup L)$ upon observing a loss. We can, however, calculate the probability of interference when there is more than one simultaneous transmission as follows:

$$P_i = P[I|S] = P[(I \cup L)|S] - P[L|S] + P[(I \cap L)|S].$$

We can calculate this conditional probability based upon events measured in the trace. For a given (s, r) pair, let n be the number of transmissions from s to r, $n_0 \leq n$ be the number of transmissions from s to r *without* a simultaneous transmission from another node, and n_0^l be the number of n_0 transmissions that are lost. Likewise, let n_x be the number of transmissions from s to r *with* a simultaneous transmission, and n_x^l be the number of n_x transmissions lost.

Then we can measure $P[(I \cup L)|S]$ empirically as n_x^l/n_x. Observing that L is independent of S, the case of simultaneous transmissions, we have $P[L|S] = P[L|\neg S] = n_0^l/n_0$ and $P[(I \cap L)|S] = P[I|S] \cdot P[L]$. A bit of algebra then reveals:

$$P_i = P[I|S] = [(n_x^l/n_x) - (n_0^l/n_0)]/(1 - n_0^l/n_0).$$

Given P_i, we can then estimate the expected number of losses during simultaneous transmissions between an (s, r) pair that are due to interference. Examining all transmissions between all sending and receiving pairs, we can estimate the extent to which interference occurs in our network.

We restrict our analysis to (s, r) pairs that exchange at least 100 packets to provide confidence in our statistical estimates. These (s, r) pairs comprise 82% of all (s, r) pairs in the trace. All such pairs experience losses with at least one simultaneous transmission. Normalizing these losses according to the background loss rate of each pair according to the above formula, we estimate that 88% of these (s, r) pairs experience loss due to interference from another node. Whose transmissions are being interfered with? Of those (s, r) pairs experiencing interference, the sender s is split roughly equally between APs (56%) and clients (44%).
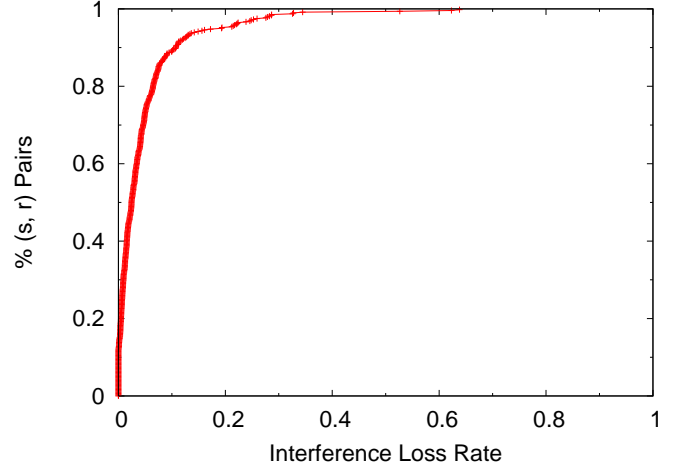

Figure 7: Interference loss rate across (s, r) pairs.

Does interference have a significant impact on the overall transmissions from senders to receivers? Again, note that lost transmissions may increase frame exchange times due to retransmissions, but not necessarily result in a failed frame exchange. To answer this question, Figure 7 shows the *interference loss rate* as a CDF across all (s, r) pairs. We define interference loss rate as the fraction of all transmissions (*i.e.*, regardless of whether there was a simultaneous transmission or not) from s to r that were lost due to interference; alternatively, it is the probability that a transmission from s to r is lost due to interference. As a baseline, the average background transmission loss rate is 0.12. In comparison, the results in Figure 7 show that many (s, r) pairs experience minor interference: 50% of (s, r) pairs experience an interference loss rate of 0.025 (a 2.5% probability of a transmission lost due to interference), or less. Yet a noticeable fraction of (s, r) pairs suffer considerably from interference: 10% of pairs experience an interference loss rate of at least 0.1, and 5% at least 0.2. A few (s, r) pairs experienced terrible interference with an interference loss rate higher than 0.5.

7.3 802.11g Protection Mode

Next we analyze the use of 802.11g protection mode in the network. We find that the protection policy by our APs is overly conservative, potentially reducing performance for 802.11g clients. We then take advantage of the global perspective provided by the distributed monitoring platform to estimate the number of 802.11g clients that would benefit from using a more practical 802.11g protection mode policy.

During busy periods, we found a high rate of CTS control frames in the trace. Investigating further, we determined that these are primarily CTS-to-self frames used for 802.11g protection (Section 2). Since protection mode increases delay and reduces throughput for 802.11g clients, APs should only use protection mode when any active 802.11b clients
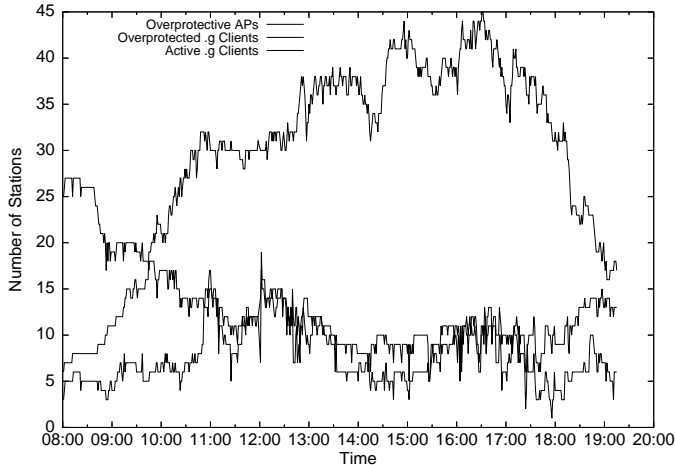
Figure 8: Overprotective APs and active 802.11g clients during the busy period of the trace.

Figure 9: TCP Loss Rate.

are in range. The APs in the network implement this protection policy, but with an overly conservative timeout. An AP will not turn off protection until an hour has passed without sensing an 802.11b client in range.

In this analysis, our goal is to identify which APs in the trace are using protection mode that unnecessarily impacts 802.11g clients; we refer to these APs as *overprotective* APs. We can identify the set of APs using protection mode based upon CTS-to-self client transmissions to those APs. Then, using the global perspective of the unified trace, for each AP using protection mode over time we can infer whether any 802.11b clients are in range of that AP after a more practical timeout of one minute. If no 802.11b clients are in range, then the AP is overprotective. We infer whether any 802.11b clients are in range of an AP using protection mode using observed probe responses. APs send these frames after they receive a corresponding probe request from a client. Our monitor density allows us to capture these responses throughout the building and create a reasonable estimate for a client's transmission range.

Figure 8 shows the impact of overprotective APs on 802.11g clients in the network for the duration of the trace. It shows (1) the total number of overprotective APs that use protection mode unnecessarily, (2) the total number of active 802.11g clients associated with these APs, and (3) the total number of active 802.11g clients in the network. During busy periods of many active clients, the number of overprotective APs decreases as more 802.11b clients become active. Similarly, the number of 802.11g clients increases and, during these busy periods, 25–50% of them are associated with overprotective APs.

A more practical protection policy would provide two benefits to clients in the network. First, the 802.11g clients associated with overprotective APs could potentially improve
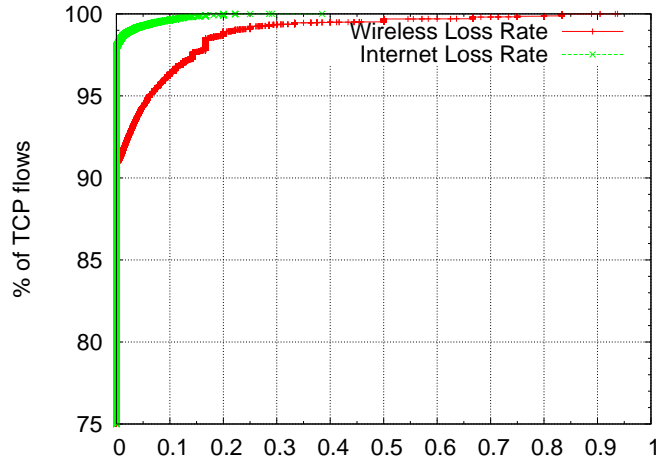
their throughput substantially. With large frames transmitted at 54 Mbps without the need for CTS-to-self, these clients could potentially improve their throughput by a factor of two.[8] Of course, this result is an upper bound: not every 802.11g client would be able to transmit at full rate, and multiple clients would still contend for the channel. However, we have found that the network is rarely at maximum utilization, even during the busiest periods. As a result, 802.11g clients should be able to benefit, especially when performing bulk transfers and the wireless network is the bottleneck hop in their path.

Second, reducing the use of CTS-to-self reduces the possibility of exposed terminals in the network, which could improve the performance of the network. Like ARP and other low-rate short frames, CTS frames have relatively high penetration and can reserve the channel across a larger space than necessary when transmitting data frames at high rates.

8 TCP Loss Rate Inference

Using the TCP reconstruction algorithm described in Section 5, we assemble all flows that complete a handshake (eliminating port scans and connection failures). From these flows we then calculate the loss rate using a variant of Jaiswal *et al.*'s approach [14]. Then, by analyzing the frame exchanges making up each TCP segment we are able to determine if each loss — as seen by TCP — is due to a lost 802.11 frame or some subsequent loss in the wired network. Figure 9 illustrates this data, showing — as expected — that the wireless component of TCP loss is dominant. What is important about this analysis is less the result itself than the capability to easily examine interactions between layers in

---

[8]CTS: 248 us (our APs send CTS at 2Mbps with the long preamble), SIFS: 16us, MSS TCP at 54Mbps: 248us, SIFS: 16us, ACK: 28us, backoff (with g): 16/2*20, backoff (with b/g): 32/2*20. The potential performance improvement is $(248 + 16 + 248 + 16 + 28 + 32/2 * 20)/(248 + 16 + 28 + 16/2 * 20) = 1.98$.

our global trace.

9 Conclusion

Network research comes to understand the artifacts it has created slowly — by careful instrumentation, monitoring and analysis. Production 802.11 wireless networks have so far escaped the level of detailed analysis experienced on the wired network – largely because of the difficulty in monitoring the wireless environment. To address this problem we have built a system called Jigsaw that unifies traces from multiple passive wireless monitors to reconstruct a global view of network activity in a production 802.11 network. We have described the algorithms used to scalably synchronize traces, unify common frames, and reconstruct the link- and transport-layer conversations embedded in those frames. Finally, we have deployed a large-scale instance of Jigsaw using over 150 monitors and used a 24-hour trace captured by our monitoring infrastructure to demonstrate complex interactions such as co-channel interference that would otherwise be difficult to analyze.

Acknowledgments

References

[1] Madwifi atheros driver for linux, 2003. `http://sourceforge.net/projects/madwifi/`.

[2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proceedings of the ACM SIGCOMM Conference*, pages 121–132, Portland, OR, Sept. 2004.

[3] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan. Characterizing User Behavior and Network Performance in a Public Wireless LAN. In *Proceedings of ACM SIGMETRICS*, Marina Del Rey, CA, June 2002.

[4] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *Proceedings of USENIX MobiSys*, San Francisco, CA, May 2003.

[5] C. Computing. 2005 National Survey of Information Technology in U.S. Higher Education, Oct. 2005.

[6] E. Daley. Enterprise LAN Grows Up, 2005.

[7] D. Duchamp and N. F. Reynolds. Measured Performance of a Wireless LAN. In *Proceedings of the 17th Conference on Local Computer Networks*, pages 494–499. IEEE, September 1992.

[8] D. Eckardt and P. Steenkiste. Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network. In *Proceedings of ACM SIGCOMM*, pages 243–254, 1996.

[9] J. Elson, L. Girod, and D. Estrin. Fine-Grained Network Time Synchronization using Reference Broadcasts. In *Proceedings of the 5th ACM/USENIX Symposium on Operating System Design and Implementation (OSDI)*, Boston, MA, Dec. 2002.

[10] Gartner. Market Share: Wireless LAN Equipment Worldwide, 2005 (Preliminary Statistics), 2005.

[11] D. Group. Wireless LAN Five Year Forecast Report, Jan. 2006.

[12] T. Henderson, D. Kotz, and I. Abyzov. The Changing Usage of a Mature Campus-wide Wireless Network. In *Proceedings of ACM Mobicom*, pages 187–201, September 2004.

[13] F. Hernández-Campos and M. Papadopouli. A Comparative Measurement Study of the Workload of Wireless Access Points in Campus Networks. In *Proceedings of the 16th International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC 2005)*, September 2005.

[14] S. Jaiswal, G. Iannaccone, C. Diot, J. Kurose, and D. Towsley. Inferring TCP Connection Characteristics from Passive Measurements.

[15] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding Congestion in IEEE 802.11b Wireless Networks. In *Proceedings of the Internet Measurement Conference*, Berkeley, CA, October 2005.

[16] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding Link-Layer Behavior in Highly Congested IEEE 802.11b Wireless Networks. In *Proceedings of the Workshop on Experimental Approaches to Wireless Network Design and Analysis (E-WIND)*, Philadelphia, PA, August 2005.

[17] R. Karp, J. Elson, D. Estrin, and S. Shenker. Optimal and Global Time Synchronization in Sensornets. Technical Report CENS-TR0012, CENS, UCLA, April 2003.

[18] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *Proceedings of ACM MobiCom*, 2002.

[19] M. McNett and G. M. Voelker. Access and mobility of wireless pda users. *Mobile Computing and Communications Review*, 9(2):40–55, 2005.

[20] A. Mishra, M. Shin, and W. Arbaugh. An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process. *ACM Computer Communications Review*, 33(2):93–102, 2003.

[21] G. T. Nguyen, R. H. Katz, B. Noble, and M. Satyanarayanan. A Trace-Based Approach For Modeling Wireless Channel Behavior. In *Winter Simulation Conference*, pages 597–604, 1996.

[22] K. N. Ramachandran, E. M. Belding-Royer, and K. C. Almeroth. DAMON: A Distributed Architecture for Monitoring Multi-hop Mobile Networks. In *IEEE International Conference on Sensor and Ad hoc Communications and Networks (SECON)*, October 2004.

[23] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan. Measurement-based Characterization of 802.11 in a Hotspot Setting. In *Proceedings of the Workshop on Experi-*

mental Approaches to Wireless Network Design and Analysis (E-WIND), Philadelphia, PA, August 2005.

[24] D. Schwab and R. Bunt. Characterising the Use of a Campus Wireless Network. In *Proceedings of IEEE Infocom*, 2004.

[25] D. Tang and M. Baker. Analysis of a Local-Area Wireless Network. In *Proceedings of ACM MobiCom*, pages 1–10, 2000.

[26] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz. Measurements of a wireless link in an industrial environment using an ieee 802.11-compliant physical layer. *IEEE Transactions on Industrial Electronics*, 43(6):1265–1282, December 2002.

[27] J. Yeo, M. Youssef, and A. Agrawala. A Framework for Wireless LAN Monitoring and its Applications. In *Proceedings of the ACM Workshop on Wireless Security (WiSe04)*, 2004.

[28] J. Yeo, M. Youssef, T. Henderson, and A. Agrawala. An Accurate Technique for Measuring the Wireless Side of Wireless Networks. In *Proceedings of the International Workshop on Wireless Traffic Measurements and Modeling*, Seattle, WA, June 2005.