Demo: Witals, AP-centric Health Diagnosis of WiFi Networks

Mukulika Maity
IIT Bombay
mukulika@cse.iitb.ac.in

Avinash Chaurasia
IIT Bombay
avinashk@cse.iitb.ac.in

Rachit Srivastava
AirTight Networks
rachit.srivastava@airtightnetworks.com

Bhaskaran Raman
IIT Bombay
br@cse.iitb.ac.in

Mythili Vutukuru
IIT Bombay
mythili@cse.iitb.ac.in

ABSTRACT

We present *Witals*, a system for WiFi performance diagnosis. In Witals, a live access point diagnoses the health of the WiFi network by examining the health of the air around it. Such diagnosis is critical functionality for sysads, and is an important addition to the state of the art. In the demo we will show WiFi performance diagnosis by Witals live at MobiCom venue and also from past collected traces in other settings & for controlled experiments.

1. INTRODUCTION & MOTIVATION

[1] Although the raw capacity of WiFi has seen tremendous growth, the real world WiFi deployments have become more troublesome for the system administrators (sysads). Diagnosing WiFi performance problem is a non trivial problem due to complex interdependence of different factors in the network stack. Witals aims to provide a sanitized view of WiFi network's health to the sysads. Prior approaches of WiFi performance diagnosis is not complete and they use heavy weight infrastructures. Witals runs at the operating AP and provides live diagnosis. We have carefully designed Witals to be low overhead, so that the AP can manage such analysis even when operating at the high data rates of 802.11n. Witals diagnoses performance problem from the perspective of downlink traffic which is a common case.

For such diagnosis, we first come up with a causal diagnosis graph which relates causes to different WiFi performance problems. Next, we identify a set of metrics corresponding to the nodes in this causal graph. Witals measures those metrics in real-time in an operational AP. The metrics quantify the impact of the performance problem. Finally using the metrics and causal relationships we design diagnosis algorithm. This ultimately presents a sanitized view of WiFi network's health to the sysad. We have implemented a prototype of Witals on an enterprise grade 802.11n AP platform. Using a variety of controlled as well as real-life measurements, we have verified that our diagnosis framework follows ground truth accurately.

[1]This work was supported in part by AirTight Networks.

MobiCom'15, September 7–11, 2015, Paris, France.
ACM. ISBN 978-1-4503-3619-2/15/09
DOI: http://dx.doi.org/10.1145/2789168.2789179.

2. CAUSAL DIAGNOSIS GRAPH

We first discuss our causal diagnosis graph shown in Fig 1. The graph relates various causes to different kinds of performance problems. The graph is designed to be complete, as it identifies all possibilities of any WiFi performance problem. The nodes in the graph are possible WiFi performance problems. The edges in the graph are causal relationships between the nodes. The underlined nodes are the various conditions Witals can diagnose. We denote AP_W to refer the access point on which Witals runs.

If the downlink throughput of AP_W is low, there are only three possible reasons. (A) Low offered load: AP_W does not have enough data to transmit, (B) Not enough airtime: AP_W has data to transmit but does not get enough opportunities to transmit, (C) Inefficient utilization of provided airtime: AP_W has data to transmit and it gets enough opportunities to transmit as well but it is not able to utilize airtime efficiently.

(A) Low offered load: marked as "A" in Fig 1. This could be caused by variety of non-wireless related reasons such as DHCP server issues, client issues etc. Witals does not distinguish between the different reasons causing low offered load.

(B) Not enough airtime: marked as "B" in Fig 1. If AP_W has data but does not get enough airtime to transmit the possible reasons are i) *Wasted non decodable energy*: The air is occupied by non decodable energy. This comprises of *non-WiFi energy* and *WiFi non-frame energy*. ii) *Control, management overheads*: Even when the energy is decodable, it could be occupied by control and management overheads. iii) *Other BSS*: Now when the airtime is taken up by data, it might belong to other BSS operating on same/nearby channel of AP_W. iv) *Upload*: Even within AP_W's BSS, airtime could be taken up by uplink transmissions. v) *MAC unfairness*: Now when above factors are not affecting AP_W's airtime, DCF MAC can result in unfairness in transmission opportunities (txop) and/or airtime. vi) Finally, even when DCF MAC is overall fair, the number of radios N on the channel might be high, although AP_W gets its fair share i.e. $\frac{1}{N}$ airtime, but due to high value of N, downlink throughput suffers.

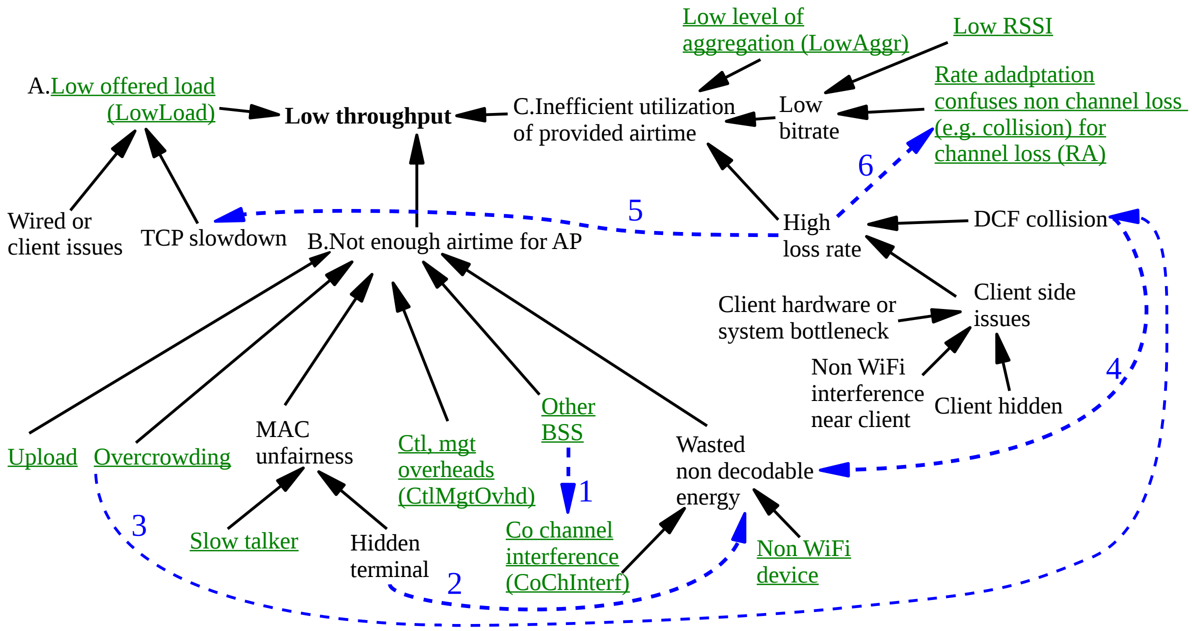

Figure 1: Witals Diagnosis: Causal Graph

(C) Inefficient utilization of provided airtime: marked as "C" in Fig 1. Finally, AP_W has data and gets enough airtime as well, but AP_W cannot utilize that efficiently. The possible reasons are i) *Low bitrate*: AP_W might be operating at a low bitrate either due to poor RSSI or rate adaptation algorithm makes wrong decisions and picks a suboptimal bitrate unnecessarily. ii) *High loss rate*: Now, even when AP_W is operating at optimal bitrate, the airtime utilization could be inefficient if link layer loss rate is high. High loss rate could be caused by DCF collision and/or client side issues. iii) *Low aggregation*: In case of 802.11n, even if bitrate is optimal and loss rate is low, still airtime utilization could be inefficient if AP_W uses lower level of frame aggregation.

While the above dependencies are from a certain perspective, for completeness, we also need several "cross-links" in the causal graph as shown in Fig 1. For instance *cross-link 1* captures the fact that if other BSS operate on same/near by channel of AP_W, some of energy might not be decodable to AP_W, this causes co channel interference.

3. WITALS FOR DIAGNOSIS

Corresponding to different nodes of the causal diagnosis graph, we come up with a suitable metric which quantifies the impact of the problem. The metrics are *comparable* with one another. We can identify and quantify multiple simultaneous causes of WiFi performance problems. We use hardware supported registers and per frame information for measurement of the metrics. We calculate these in windows of time, of duration denoted as T_W. Table 1 lists the various possible diagnoses of Witals and the corresponding metrics to diagnose those & also how we measure the metrics.

4. DIAGNOSIS ALGORITHM

Using the causal diagnosis graph and the corresponding metrics we develop the diagnosis algorithm. The algorithms gives a refined view to the sysad representing the health of the network. The algorithm runs at the end of every T_W. It gives diagnosis as one or more as listed in Table 1. In our implementation we choose $T_W = 2s$.

Diagnosis	**Metric**
Healthy	-
LowLoad	I (idle airtime)
NonWiFi	W (airtime lost due to wasted non-decodable energy)
CoChInterf	W, OB (airtime occupied by frames of other BSS)
CtlMgtOvhd	O (airtime occupied by control and management frames)
OtherBSS	OB
Upload	Ul (airtime occupied by uplink transmissions)
SlowTalker	S_{redn} (how much more airtime AP_W could have achieved if all radios were at least as fast as AP_W)
Overcrowding	W, N (number of radios)
LowRSSI	$RSSI_{redn}$ (throughput reduction for operating at lower than ideal bitrate due to poor RSSI)
RA	RA_{redn} (throughput reduction due to operating at lower rate as opposed to operating at ideal RSSI bitrate), N
LowAggr	$Aggr_{redn}$ (throughput reduction for operating at lower than ideal level of frame aggregation)

Table 1: Diagnosis Table

Algorithm is summarized in Algorithm 1. At each T_W we pick three major throughput reduction metrics call this set as $\mathbb{M}$. We apply a minimum threshold percentage (here 10%) for any metric to have an impact on performance. If $\mathbb{M}$ is empty we diagnose *Healthy*. Given $\mathbb{M}$ is nonempty, if idle percentage is more than a threshold (here 50%), i.e. mostly idle, then we diagnose *LowLoad* and *Healthy*. The diagnosis of *CtlMgtOvhd*, *OtherBSS*, *Upload*, *SlowTalker*, *LowRSSI*,

Algorithm 1 Diagnosis algo, runs every T_W

$\mathbb{M}$ = Set of up to 3 major witals metrics > (10%)
if $\mathbb{M}$ *is* $\emptyset$ **output** Healthy
if $I > 50$ **then** ▷ Channel more idle than busy
  output LowLoad, Healthy
else ▷ Channel more busy than idle
  if $O \in \mathbb{M}$ **output** CtlMgtOvhd
  if $OB \in \mathbb{M}$ **output** OtherBSS
  if $Ul \in \mathbb{M}$ **output** Upload
  if $S_{redn} \in \mathbb{M}$ **output** SlowTalker
  if $RSSI_{redn} \in \mathbb{M}$ **output** LowRSSI
  if $RA_{redn} \in \mathbb{M}$ **output** RA
  if $Aggr_{redn} \in \mathbb{M}$ **output** LowAggr
  if $W \in \mathbb{M}$ **then**
    if $N > 5$ **output** Overcrowding
    if $OB > 10$ **output** CoChInterf
    if neither of the above **output** NonWiFi
  end if
end if

RA, *LowAggr* is straightforward, we simply check if the corresponding metric is in $\mathbb{M}$. Some problems are trickier to diagnose because their impact cannot be captured in any one metric. Overcrowding, non WiFi device and co-channel interference all cause increase of W. So, we use other metrics to distinguish between these. If N is more than a threshold (here 5, empirically observed) we diagnose *Overcrowding*. The presence of significant interference from other BSS is used to flag *CoChInterf*. But if neither of the above is true we conclude diagnosis as *NonWiFi* device.

Diagnosing at every T_W might be too fine grained as many transient problems could be shown up. So if the sysad is more interested in diagnosing persistent performance problems, the framework allows to run the algorithm for long term by taking the output of previous T_W algorithm over multiple such windows as input.

5. WITALS DEMO

At the demo we will show the performance diagnosis by Witals for controlled experiments and for the WiFi deployment at MobiCom venue & also from past collected traces in other settings. While Witals is designed to run on an operational AP, it can equivalently run from a nearby sensor in the same channel as the AP. Also note that the diagnosis is passive: no extra traffic is generated.

Controlled experiments: For controlled experiments, we will set up our access point in 5Ghz band, so that it does not interfere with existing WiFi network there. We will perform experiment for diagnosis of *OtherBSS* , *SlowTalker* and *LowRSSI* by artificially creating each problem.

Live demo: We will also show live performance diagnosis at the venue. Witals will diagnose presence of any performance problem in the network. Thus we can examine the health of the WiFi network at the venue and diagnose the persistent problems (if any) for better deployment of WiFi for such conferences in future.

Now to better understand what the demo will show, we illustrate using examples. From our real life experiments we take two real life scenarios, classroom and conference.

Classroom: We ran Witals in a classroom where 80 clients were associated with AP_W. The class activity was

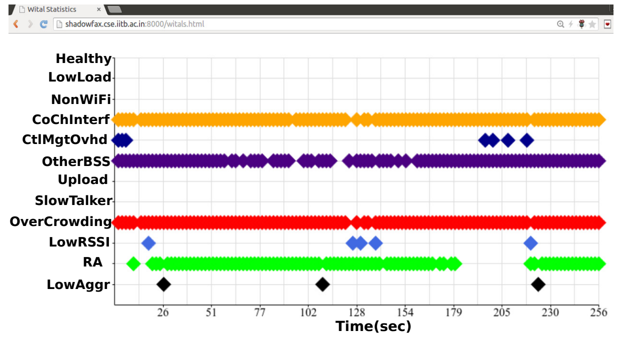

Figure 2: Diagnosis screenshot: classroom

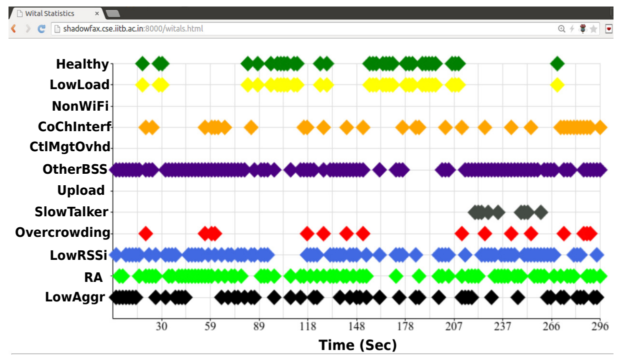

Figure 3: Diagnosis screenshot: conference

online quiz using tablet devices. Fig. 2 shows Witals' diagnosis plotted over time. We can see here that *Overcrowding*, *OtherBSS* and *CoChInterf* are the major performance problems for this network.

Conference: We ran Witals at a research conference. Here AP_W acted like a sniffer, collected metrics of operating AP. We did analysis from the saved metrics. There were many access points on same and nearby channel of operating AP. Fig. 3 shows Witals' diagnosis plotted over time. Unlike classroom, *Overcrowding* is not a major problem here. This is expected, as unlike classroom users do not use the network simultaneously. In several time windows the network is healthy and lightly loaded. As pointed out by diagnosis, *OtherBSS*, *LowRSSI*, *RA* and *LowAggr* are major performance problems in this network. *LowRSSI* problem might have been caused as the conference venue was across two floors.

6. SETUP

Equipment: we will bring our own Witals enabled APs (2-3).

Space needed: one table is sufficient, set up time is around 5min.

Addition facilities: power is needed, Internet is not required.

7. ACKNOWLEDGEMENT

This work was supported in part by AirTight Networks.