Proceedings of the 2015 International Mechanical Engineering Congress & Exposition
IMECE 2015
November 13-19, 2015, Houston, Texas, USA

IMECE2015-50998

THE APPLICATION OF AXIOMATIC DESIGN TO COMPLEX WI-FI SYSTEMS

Jason D. Hintersteiner, CWNA, CWDP, CWAP, CWSP
President & Chief Technology Officer
Imperial Network Solutions LLC
P.O. Box 353, Wilton, CT 06897 USA

ABSTRACT

Axiomatic Design is a technique that has been applied to multiple disciplines for enabling design, analysis, and troubleshooting of complex systems. In this paper, the principles of Axiomatic Design are applied to Wi-Fi networking. Wi-Fi is the information backbone for numerous applications, including Internet connectivity, video surveillance, data collection and inventory tracking in manufacturing and warehouse environments, patient location and health status monitoring in assisted living and hospital environments, along with numerous others. A Wi-Fi system consists of multiple access points working in tandem to provide seamless, high-speed, and high-quality wireless coverage to one or more wireless client devices. To implement such a network effectively, the Wi-Fi engineer must understand and control the interactions between multiple engineering disciplines, most notably information technology, network engineering, radio frequency physics, antenna design, and materials science. Technology development in this field is fast-paced, with new standards and capabilities being introduced into the market every couple of years. Additionally, the customer expectations (i.e. requirements) are changing as well once the Wi-Fi network is installed, as data demands from new types of devices like smartphones, tablets, and network appliances are introduced long after the original network was implemented.

This paper shows that there are three primary functional requirements for a Wi-Fi network, namely client usage type, coverage area, and client capacity. When designing, implementing, or troubleshooting a Wi-Fi network, there are four primary design parameters that can be controlled, namely AP antenna / model, location, channel, and transmission power. Axiomatic Design demonstrates that these four design parameters are coupled, and thus cannot be manipulated independently. Nevertheless, by effectively implementing Axiomatic Design techniques to define a set of *best practices*, these four key parameters can be decoupled and properly linked back to the requirements and constraints of the system to simplify the design, implementation, and troubleshooting of a Wi-Fi network.

INTRODUCTION

By the end of 2014, there were more mobile-connected devices than people on Earth. By 2019, it is projected that there will be 26 billion mobile devices, or more than 3 mobile devices per capita, with the overwhelming majority either fully or partially depending on Wi-Fi networks for maintaining Internet connectivity. [7,8] Primarily, this growth has been driven by devices capable of consuming more data and streaming HD and 4K video to small screen devices, such as smartphones and tablets. Additionally, there is an emerging trend, known as the Internet of Things (IoT), where all manner of physical objects are being embedded with electronics, software, sensors, and wireless connectivity to add value by communicating data back to the device manufacturers and/or cloud-based applications. While many of the IoT projections may prove to be marketing and analyst hype, some of these devices are already available on the market and being adopted. Accordingly, the demand for constant and ubiquitous wireless connectivity is already here, and will continue to grow exponentially. [8]

The IEEE 802.11 standard, commonly known as Wi-Fi, was originally launched in 1997, with several amendments released every few years to improve bandwidth, security, power management, and quality of service (i.e. traffic prioritization). Over the past decade, Wi-Fi deployments have grown exponentially, in both consumer and enterprise markets, to attempt to quench our society's insatiable thirst for more bandwidth and ubiquitous connectivity.

Despite the common misconception that a Wi-Fi deployment is as simple as purchasing a "black box" at a consumer-goods electronic store for $50 and plugging it in to get Wi-Fi, the reality is far more complex. Wi-Fi is a radio frequency (RF) modulation technology operating in unlicensed spectrum. As such, Wi-Fi is subject both to the laws of RF

physics and to interference from other neighboring Wi-Fi and non-Wi-Fi sources. Deploying Wi-Fi in complex environments, such as schools, universities, hospitals, hotels, apartment buildings, shopping malls, retail outlets, and even large private houses, multiple access points must be used. Great care must be taken to ensure that these access points are working collectively and not at odds with each other, making the network of access points (APs), switches, and routers a complex engineering system in its own right.

Like any complex engineering system, the design of a Wi-Fi system is both driven by and measured against its requirements and constraints. [13] Accordingly, requirement and constraint gathering is a critical step, unfortunately quite frequently skipped, at the start of any design or troubleshooting effort. The type of equipment chosen, where it is placed, and how it is configured are highly dependent upon what the customer needs the system to do and what physical and logical constraints must be worked around. Each access point equipment vendor tends to have a particular niche for which their equipment is well-suited, and particular requirements will often drive the selection of particular vendors and equipment models.

Accordingly, Wi-Fi networks are a true example of *mass customization*, a term used to describe manufacturing systems capable of leveraging mass-production techniques for cost efficiencies while still personalizing the products at individual or low volumes for particular applications and needs. Wi-Fi networks do tend to have similar requirements both within and across vertical markets. Nonetheless, a Wi-Fi system must always be tailored to both the physical layout and building materials of the structure in which it is installed, as well as to the number and types of devices as well as how those devices are going to use the network for wireless connectivity. It is clearly not cost-effective to "re-invent the wheel" for every deployment, so it is generally convenient for Wi-Fi engineers and installers to have a small number of pre-determined sets of network equipment, configuration standards, and installation practices to deploy such systems for customers.

The danger, of course, is that many installers tend to deploy either the access point technology they know, or conversely the access point technology they can purchase at the cheapest price on a particular week. The reality is that particular vendors tend to focus on particular verticals, and thus be better (or worse) suited to particular types of applications and deployments. Years ago, when Wi-Fi requirements were not too stringent (i.e. "best effort"), even a poor quality deployment often was "good enough" to meet the customer's needs. Nowadays, the requirements are not only more stringent because of the desire to push the limits of performance, but these requirements change over time as the performance demands continue to increase. Thus, a deployment that was considered acceptable even just a couple of years ago may no longer be viable, and an installation done today may not be considered viable in a 3-4 year timeframe.

As access point technology gets increasingly more complicated, and as the required demands on a Wi-Fi network get ever more difficult to provide higher data capacity and higher connection reliability, design tools are required to ensure that the Wi-Fi systems are properly and optimally architected to meet their functional requirements. In this paper, the system design principles based on *Axiomatic Design* are applied to the design of Wi-Fi systems. To the author's knowledge, this paper represents the first attempt ever to bridge a system design methodology from the Mechanical Engineering / Systems Engineering discipline and apply it to the field of Wi-Fi networking.

AXIOMATIC DESIGN FOR SYSTEMS

Axiomatic Design provides a framework for describing "design objects" which is consistent for all types of design problems and at all levels of detail. Thus, different designers can quickly understand the relationships between the intended functions of an object and the means by which they are achieved. Additionally, the design axioms provide a rational means for evaluating the quality of proposed designs, and the design process that is used guides designers to consider alternatives at all levels of detail and to make choices between these alternatives more explicit.

In Axiomatic Design terminology, design is defined as the development and selection of means (DPs) to satisfy objectives (FRs), subject to constraints (Cs). The main concepts of Axiomatic Design include: (1) *domains*, which separate conceptually the functional and physical parts of the design; (2) *hierarchies*, which categorize the progress of a design in the functional and physical domains from a system level to more detailed levels; (3) *zigzagging*, which indicates that the decisions made at one level of the hierarchy affect the problem statements at lower levels; and (4) *design axioms*, which dictate that the independence of the FRs must be maintained (Independence Axiom) and that the information content (i.e. cost, complexity, etc.) must be minimized (Information Axiom) in order to generate a design of good quality. [1,2]

By definition, functional requirements (FRs) are independent of each other. Ideally, each functional requirement should have one, and only one, corresponding design parameter (DP), and that design parameter should only influence its corresponding FR. This case is known as an *uncoupled design*. As designs get more complex and more constrained, this is usually difficult, if not impossible, to achieve in practice. A *coupled design* is the case where all of the DPs impact all of the FRs. This is the situation that needs to be avoided, as the change to any single FR, and thus to its corresponding DP, impacts all of the other FRs, thus requiring changes to other DPs to compensate. Such a design requires iteration and thus becomes very difficult to optimize, as well as being very fragile in responding to even minor requirements changes. In these cases, DPs should be chosen and/or limited to provide a *decoupled design*, so that the DPs can be changed in a

particular sequence so as to affect the FRs without requiring further iteration.

It was noted above that, for an *ideal design*, the number of FRs equals the number of DPs. The case where there are more FRs than DPs is known as *insufficient*, as it is impossible to independently satisfy all of the FRs, because there is an insufficient number of DPs to do so. The case where there are more DPs than FRs is known as *redundant*. In this case, one or more of the additional DPs can generally be held fixed, so as to allow independent or sequential manipulation of the other DPs. [1]

The classic demonstration of the principles of Axiomatic Design is a water faucet. There are two FRs: (1) control the water flow rate, and (2) control the water temperature. Since water is generally supplied to a faucet via two pipes, one transporting hot water and the other cold water, the simplest solution, as shown in Figure 1, is to have two faucets, one for the hot water flow and one for the cold water flow. These valves are the DPs. Anyone who has ever used such a faucet knows that this is a *coupled design* – you cannot adjust either valve without affecting both flow rate and temperature.

$$\begin{bmatrix}\text{FR1: Flowrate}\\ \text{FR2: Temperature}\end{bmatrix} = \begin{bmatrix}X & X\\ X & X\end{bmatrix}\begin{bmatrix}\text{DP1: Hot Faucet}\\ \text{DP2: Cold Faucet}\end{bmatrix}$$

Figure 1: Coupled water faucet design.

Contrast this with a faucet with a mixer tap valve, as shown in Figure 2, such that moving the lever vertically controls the flow rate by drawing in water from both pipes equally, while moving the lever horizontally controls the temperature by deliberately changing the valve size (i.e. inlet areas) of each pipe, thus altering the ratio of hot water to cold water. Both parameters can now be controlled independently, and thus is an example of an *uncoupled design*.

$$\begin{bmatrix}\text{FR1: Flowrate}\\ \text{FR2: Temperature}\end{bmatrix} = \begin{bmatrix}X & O\\ O & X\end{bmatrix}\begin{bmatrix}\text{DP1: Mixer Valve Vert.}\\ \text{DP2: Mixer Valve Horiz.}\end{bmatrix}$$

Figure 2: Uncoupled water faucet design.

The mixer tap valve imposes the constraint that the pressure and size of the hot and cold water feeds, and thus their volumetric flow rates, are equivalent. What would happen, however, if one of the feeds had a disproportionately higher flow rate? In this scenario, adjusting the vertical lever of the mixer tap valve would impact both the flow rate and the temperature. However, adjusting the horizontal lever of the mixer tap valve still only impacts the temperature (i.e. the ratio of hot to cold water), and not the flow rate. By adjusting these two parameters in sequence, i.e. first the vertical lever for flow rate and then the horizontal lever for temperature, both parameters can be adjusted without further iteration. This is an example of a *decoupled design*.

$$\begin{bmatrix}\text{FR1: Flowrate}\\ \text{FR2: Temperature}\end{bmatrix} = \begin{bmatrix}X & O\\ X & X\end{bmatrix}\begin{bmatrix}\text{DP1: Mixer Valve Vert.}\\ \text{DP2: Mixer Valve Horiz.}\end{bmatrix}$$

Figure 3: Decoupled water faucet design.

A *system design* is defined as a design in which the task of the designers is to integrate several DPs into a whole in order to satisfy some set of FRs, subject to a given set of constraints. This differs from other design methodologies in that the process followed is: (a) identify the FRs and their corresponding constraints, (b) determine possible DPs, and (c) integrate them into a system. In this case, an understanding of the interrelationships among the different DPs is vitally important. By not doing so, the design process becomes a confusing muddle, which can ultimately lead to several iterations and poor design decisions. In system design, it is desirable that the same process be useful at all levels of the design; the same approach may be followed recursively, starting at the system level and continuing until the design is complete. Thus, the system architecture is useful both when design concepts are initially generated and evaluated, as in new designs, and when changes are proposed to existing designs. [3]

By comprehensively documenting the interrelationships between the FRs, Cs, and DPs at every level of the design hierarchy, a *system architecture* can be generated and used as a tool for decision making. The strength of the system architecture is that, in addition to the operational flow of the system, it also captures the order in which design decisions have to be made, and indicates how the alteration of one part of the system can potentially impact other parts. [4]

For large systems, the system architecture breaks down the design into individual systems and subsystems at each level of the design hierarchy. In this representation, a system is modeled as a series of interacting inputs and outputs, and its functions are broken down into three categories: process functions (i.e., functions that perform value-added activities), command and control logic (e.g. software, controllers, etc.), and support and integration functions (e.g., pneumatics, mechanical structure, wired network backbone, etc.). [5,6]

WI-FI NETWORKS: FUNCTIONAL REQUIREMENTS AND CONSTRAINTS

While each Wi-Fi system deployment has a unique set of requirements and constraints, these can be generally grouped into a few basic categories, which define our top-level functional requirements (FRs) and constraints (Cs).

FR1: Usage

One of the key requirements in defining a Wi-Fi system is to identify the types of devices that will be using the network, the types of applications to be run on those devices, and the level of security that needs to be provided. This is naturally unique to each environment, but tends to cluster around the needs of particular vertical markets. As an example, a hotel will have a different set of users and devices on a nightly basis, but the types of devices brought in (e.g. smartphones, laptops) and the need to separate the network for guests from the network for hotel staff (e.g. maintenance, housekeeping, front desk, etc.) are common across different hotel properties.

When more than one type of network is needed in a facility, which is increasingly common even in small deployments, each wireless service is provided with its own SSID and placed on its own virtual local area network (VLAN), which allows the division of one physical Wi-Fi system into multiple parallel virtual systems that are normally isolated from each other, as illustrated in Figure 4.

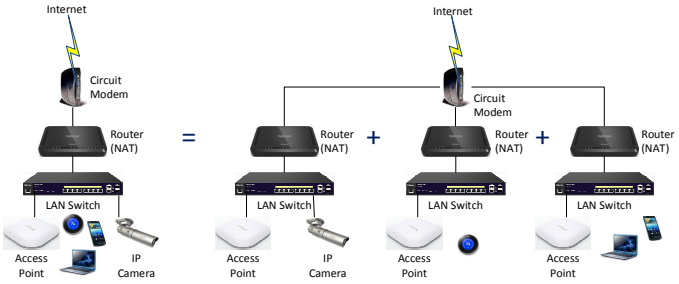

Figure 4: Concept of Virtual Local Area Networks (VLANs). One set of hardware is used to create multiple parallel virtual networks, each with their own unique set of usage requirements.

Generically, most enterprise Wi-Fi environments will have some or all of the following types of network usage requirements.

- **Visitor / Guest Network:** This network is for people visiting the facility on a temporary basis, generally with smartphones, tablets, and /or laptops, and typically require access to the Internet but not access to any other resources on the internal network. In addition to isolating users from internal network resources, users should also be isolated from each other (i.e. not see each other on the local network). Encryption of the wireless signal is generally not used, so as to facilitate ease of access to the network.
- **Staff Network:** This network is for the staff of the facility, and is intended for devices necessary for operations, which can includes tablets, laptops, PCs, and/or dedicated network appliances. This network is typically encrypted to prevent unauthorized access, utilizing either pre-shared key security or access control via a RADIUS or Active Directory Server.
- **Voice Network:** For facilities using Voice over IP and/or Voice over Wi-Fi devices (VoIP or VoWiFi), a dedicated SSID is established with network traffic prioritization. Wireless traffic is usually encrypted with a pre-shared key to facilitate roaming between access points.
- **Security Network:** This network utilizes both appliances (e.g. IP cameras, access control locks, etc.) which may be connected wired or wirelessly, along with monitoring from fixed and/or portable security stations. Wireless traffic is usually encrypted with a pre-shared key to prevent unauthorized access.
- **Network Appliances:** As IoT becomes more commonplace, enterprises start using appliances for improving operational efficiencies. This network generally contains appliances for lighting and temperature control (e.g. NEST thermostats), along with multimedia and other "smart home" appliances. These appliances may not individually be passing a lot of data, but there could potentially be quite a lot of them on the network, depending on the size of the facility.

FR2: Coverage

There is a requirement to provide coverage in all of the locations of a facility where Wi-Fi devices and applications are to be used. This requirement may seem straightforward, but is influenced heavily by two types of constraints:

- **Physical Constraints:** These are the physical attributes of the facility, including the layout and the (lack of) availability of cabling paths to connect access points. If the facility is new, wiring can often be run inexpensively before the walls are fully erected and sealed. In an existing facility, running new wiring can often be extremely challenging and/or expensive. Even more importantly, however, are the building materials that make up the walls, ceilings, and floors, because radio signals pass much more easily through certain types of materials than others. As an example, open air cubicles, uncoated glass windows, and even drywall are fairly transparent to RF, whereas brick or concrete will highly attenuate Wi-Fi transmissions. There are some materials, like wire-mesh stucco, low-e glass commonly used for outer windows of LEED-certified buildings, and reinforced concrete that serve to block virtually all RF penetration.
- **Logical Constraints:** These are the constraints imposed by owners and operators of the facility. The largest of these is generally budget, though aesthetics, especially in high end properties, can be a major factor. There may also be other systems that have to be worked around, such as a legacy or neighboring Wi-Fi or unrelated radio systems using the same frequencies.

FR3: Capacity

In addition to the types of devices using the network, the number of simultaneous devices and the average and peak traffic loads expected by those devices is a critical factor to designing a Wi-Fi network. In high capacity areas, such as classrooms, conference centers, arenas, or stadiums, one AP may be more than sufficient to provide coverage, but multiple APs could be needed to handle the high number and traffic volume of devices. Accordingly, the area may need to be covered in smaller "cells" in order to handle the quantity of devices and the traffic volume they generate.

Within this capacity requirement is a projection of how capacity needs are going to grow over time. Large enterprise Wi-Fi deployments have a typical expected lifespan of 3-5 years, and smaller enterprise deployments have an expected lifespan of 5-7 years. Meanwhile, the requirements for both bandwidth consumption and quantity of devices tends to more closely follow Moore's Law (i.e. doubling every 18 months). Accordingly, the network installed today may be perfectly adequate and have margin for today's need, but will be inadequate in perhaps only 2-3 years. Hence, when selecting components and designing for capacity, it is tomorrow's loosely defined, requirements, not today's, which drive the design.

FR4: Control

In keeping with the use of Axiomatic Design for complex systems, a function is needed to coordinate and control the Wi-Fi network. Common requirements in this category are as follows:

- **Authentication / Association:** A client device is required to establish and maintain a connection to the network.
- **Access Control:** Generally, there is a need to control who (i.e. what devices) are allowed to access the network. When there are multiple VLANs in place, these user types must be distinguished and placed on the appropriate virtual local area network.
- **Security / Encryption:** Many applications require the use of encryption across the wireless link to prevent unauthorized access to data.
- **Roaming:** When devices move from one area of the facility to another, they will eventually lose contact with the original access point. Hence, there are typically requirements for client devices to move around the facility and roam between access points without a noticeable impact on application performance.
- **Monitoring and Maintenance:** Most Wi-Fi systems need to be monitored to ensure that all components are online and functioning properly. In case of failure events, some systems may need automatic processes enabled to attempt to resolve the issue or fail into a reduced, but still functional, state.

FR5: Support & Integration

In keeping with the use of Axiomatic Design for complex systems, a support and integration function is required to tie the various functions of the Wi-Fi network together. Common requirements in this category are as follows:

- Provide power to access points
- Provide data connectivity to access points
- Provide infrastructure to facilitate communication between the access points and the central router, central controller, and/or the Internet.
- Provide enclosures to protect access points from physical tampering and/or environmental effects (e.g. harsh temperature or weather conditions)

A Note on FR Independence

It can be debated whether capacity {FR3} is truly independent from usage type {FR1} and coverage area {FR2}, and thus meet the definition of FRs as the minimum set of independent requirements that completely characterizes the functional needs of the system [1,2]. Strictly speaking, there are some implied dependencies between these parameters that are represented as intrinsic constraints. As in any complex system, these intrinsic constraints can effectively restrict your DP selection to a degree where coupling is unavoidable. Nonetheless, in the strictest sense, these FRs are independent, and, as will be shown later, their corresponding DPs can be selected so as to decouple the design.

The independence of these FRs is best illustrated by example: Take two hotels of the same commercial chain, constructed to the same blueprints and specifications, but geographically located in two different cities, namely Honolulu, HI and Cambridge, MA. Both hotels require a guest network for hotel guests and a staff network for hotel employee operations {FR1}. Additionally, both hotels have the same layout and building materials, thus have identical areas requiring coverage {FR2}. The capacity requirements of the networks in these two hotels {FR3}, however, are quite different because of the clientele each hotel caters to, and thus how those hotel guests use the network. The guests at the hotel in Honolulu are primarily vacationers. While vacationers do bring several network devices with them (e.g. tablets, smartphones, laptops, gaming consoles, e-readers, etc.), most of the activity is spent enjoying the tropical climate, so such devices are mostly unused or are only used for brief periods, with the exception being the outdoor pool area, where lots of e-readers are in use every day while vacationers sunbathe. Cambridge, MA, however, is a hub for high tech businesses, and the hotel in this city caters to business travelers, who are bringing more devices (i.e. each average guest likely has a smartphone, tablet, and laptop), using more devices simultaneously and consuming significantly more bandwidth as they spend time during the evening hours working on presentations and other business-related tasks. Furthermore, the

outdoor pool is closed much of the year due to the climate of Cambridge, MA, and only gets light usage during the summer months. The hotel in Cambridge, MA is, on average, going to have more simultaneous devices consuming more bandwidth than the hotel in Honolulu, HI, thus necessitating a larger bandwidth pipe and, potentially, more access points and switches to handle the increased number of devices.

WI-FI NETWORKS – DESIGN PARAMETERS

When assembling a Wi-Fi design, there are four fundamental "knobs" that can be turned in the design, making up the design parameters.

DP1: Access Point / Antenna Type

There are multiple access point (AP) vendors providing products with multiple types of capabilities. Many vendors focus their product lines to meet particular types of challenging environments (e.g. stadiums / arenas, K-12 education, etc.), and are optimized and priced accordingly. The most critical, and often most challenging, design decision is the decision of vendor and product line. While not comprehensive, the following list characterizes the major differences between most product vendors and models. [9,10,11]

- **Wi-Fi generation:** The IEEE has extended the original 1997 definition of Wi-Fi numerous times, sometimes with small changes, and other times with entire new product generations. At the time of this writing, 802.11ac is the current product generation, with "wave 2" products beginning to supplant "wave 1" products introduced in 2014. Most vendors also provide models that are of the previous generation (e.g. 802.11n) at lower price points. The most modern APs on the market today will likely be 2-3 generations behind in 4-5 years. Newer generations allow for more complex modulation and coding schemes, which enable faster data rates.
- **Frequency band(s):** 802.11n allowed the use of both the 2.4 GHz frequency band and the 5 GHz frequency band. 802.11ac works entirely on the 5 GHz frequency band, though most access points of this generation will incorporate dual band capability, with 802.11n support on the 2.4 GHz band. Some single band access points are manufactured for particular applications, such as point-to-(multi)point wireless links.
- **MIMO:** 802.11n introduced MIMO (multi-in / multi-out) to allow for parallel data transmissions between multiple transmit and receive antennas. More radios allow for higher throughput, at the expense of power consumption and space. 802.11ac wave 2 shall introduce multi-user MIMO (MU-MIMO), allowing the AP to talk to simultaneously communicate with multiple clients in the same environment. MIMO and MU-MIMO allow for denser utilization of the channel, thus increasing total potential channel capacity and throughput per client device.
- **External vs. Internal Antennas:** From an aesthetic standpoint, most people don't want to see antennas, so built-in antennas are quite popular. External antennas, however, offer additional flexibility for mounting, along with the ability to add 3rd party antennas with particular profiles for custom applications.

DP2: Locations

This DP is the number and layout of the access points within the facility. While seemingly straightforward, there are many factors to consider. First and foremost, where can APs be placed such that cabling can reach it to provide power and data backhaul. Aesthetic constraints may also limit where APs can be placed within a facility. Placing APs near, or on, metal objects such as pipes, ductwork, I-beams, etc. can change its coverage profile, as those materials serve to act as antennas and distort the coverage profile. Furthermore, APs with internal antennas are generally designed for a particular orientation (e.g. a horizontal ceiling mount), and thus mounting them differently (e.g. vertically on the wall) will change the coverage profile. Finally, APs that are in line of sight of each other (e.g. APs mounted down a long hallway) are more likely to interfere with each other than if there is physical structure between the APs that will serve to partially attenuate the signals between the APs (e.g. APs mounted every few rooms on alternate sides of the hallway). [9,10,11]

DP3: Channel

Each access point broadcasts a signal on a particular channel, which is specified as a particular center frequency and channel width. On the 2.4 GHz band (802.11b/g/n) in North America, there are 11 channels of 20 MHz size allowed by the FCC. (Channels 12-14 are allowed in some other countries, such as Japan). However, the center frequencies of channels 1-13 are only 5 MHz apart, leading to only three non-overlapping channels, as shown in Figure 5.

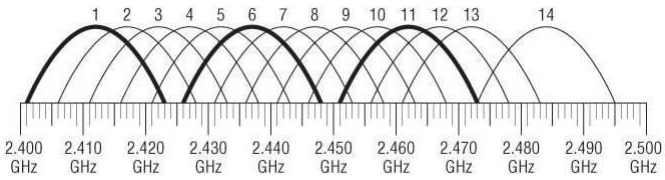

Figure 5: 20 MHz channels on the 2.4 GHz frequency band.

The 802.11n spec allows for the optional use of 40 MHz channels on the 2.4 GHz band, by bonding two neighboring channels together. However, given that the entire usable band is only 72 MHz wide, there are no two 40 MHz channel sizes that are independent, as shown in Figure 6. This makes the use of 40 MHz channels impractical in multi-AP deployments, though it is still unfortunately fairly common to see in practice as most vendors allow this channel width in their default settings.

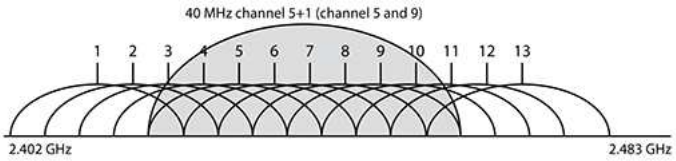

Figure 6: 40 MHz channels on the 2.4 GHz frequency band.

The 5 GHz band is much larger (over 555 MHz, semi-contiguous), and thus makes selecting independent channels and using larger channel widths via bonding neighboring channels much simpler. 802.11a allowed the use of 20 MHz channels. 802.11n allows the use of 40 MHz channels, and 802.11ac allows the use of up to 80 MHz or 160 MHz channels. This is shown in Figure 7. [9,10,11]

Note that over 2/3 of the frequency space, however, is also used by legacy military, radar, and weather systems, leading to FCC requirements to detect and move off those channels if such external systems are detected. As a result, much of the UNII-2 and UNII-2e bands are not supported by some consumer devices, leading to only two 80 MHz channels and zero 160 MHz channels.

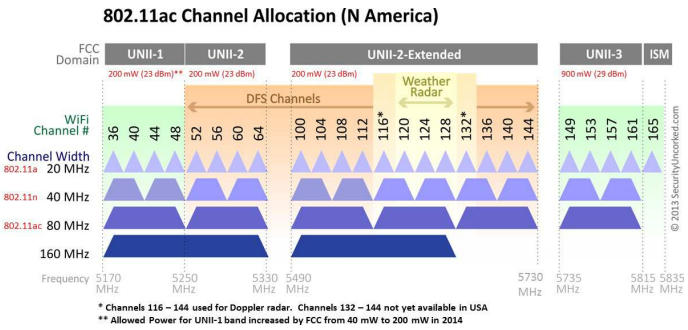

Figure 7: Channels on the 5 GHz frequency band.

DP4: Transmit Power

Transmit power of a radio is proportional to its effective range – the higher the transmit power, the more distance that a signal can travel, and/or the more physical materials that it can penetrate, and still be resolved at the receiver. Additionally, a stronger signal at a given distance generally results in a higher signal to noise ratio, allowing for more complex modulation and coding schemes and thus faster data speeds.

In early Wi-Fi deployments, which were primarily driven by the FR for coverage, it was common to turn up the power on the AP transmitter as high as is allowed by FCC and IEEE regulations. This approach was sufficient when most clients had reasonably strong transmitters themselves, such as laptops.

With the emergence of smartphones, tablets, and network appliances, however, there is often a transmit power mismatch that leads to a range mismatch. Most smartphone, tablet, and appliances use relatively weak transmitters in order to preserve both space and battery life. As a result, the situation develops where the client device can receive the relatively strong transmissions of the access point, but the access point cannot receive the relatively weak transmissions of the client device in response. Accordingly, though non-intuitively, the effective coverage area is driven more by the client devices, and the AP power levels should be set to better match the limitations of the clients.

Finally, as compared to 5 GHz, 2.4 GHz has less free space path loss and attenuation through standard building materials, giving it a larger effective range at a given transmit power level. When using a dual band access point, one generally wants to have the coverage area equivalent for both bands. This generally leads to a 4-6 dB difference in power levels on the 2.4 GHz band as compared to the 5 GHz band. In high density environments, it is not unusual to install a denser deployment of APs and then disable the 2.4 GHz band on some of them. [9,10,11]

DP5: Network Management System

Wi-Fi network activities need to be controlled, coordinated, and monitored. Many access point vendors use access point controllers with relatively "thin client" APs, so that the intelligence of the network is coordinated by a central appliance. Other vendors use standalone APs (i.e. "thick client" APs) where the APs coordinate directly amongst themselves, using a network management system (NMS) to collect usage statistics and log data.

There are three types of AP controller architectures that are commonly implemented:

- **Central Architecture:** In this scheme, all of the intelligence of the network is at the AP controller appliance on the network, and all traffic from the access point is tunneled to the AP controller before being routed to the appropriate destination. As Wi-Fi speeds have increased, the AP controller can become the bottleneck for performance, so this approach is generally no longer used.
- **Distributed Architecture:** In this scheme, all of the intelligence of the network is at the APs themselves, and an AP controller may not even be installed on the network, or if it is, only serves to collect usage statistics and coordinate AP configuration and firmware upgrades. This approach can prove problematic in more complex environments, due to the difficulties in coordinating operational functions across APs, such as client device roaming.
- **Split Architecture:** In this scheme, the intelligence of the network is split between the AP controller and the individual APs. The implementation of this varies by vendor, though typically all data handling functions would be handled by the AP, while management and control functions are handled by the AP controller.

It is also common for wireless networks to be monitored and managed remotely from a remote location, such as a centralized network operations center (NOC). Many vendors

have also introduced "cloud controllers", which are AP controllers that are located on a hosted server on the Internet, managing multiple individual network locations, each consisting of multiple APs.

DP6: Wired Network

Fundamentally, an access point is a device that allows one or more wireless client devices to connect to a wired network. The wired network supporting the wireless access points is, in and of itself, a complex system that requires many components, such as cabling, switches, routers, and modems. The application, coverage, capacity, and control FRs drive the need to provide a low-voltage cabling and switch infrastructure that meets these requirements and does not itself become a bottleneck in communications.

DECOUPLING A WI-FI NETWORK: BEST PRACTICES

In the most general sense, a Wi-Fi network based on the FRs and DPs discussed in the previous section is a coupled design. Since there are more DPs than FRs, this is also a redundant design. This can be seen in Figure 8.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3: Capacity} \\ \text{FR4: Control} \\ \text{FR5: Integration} \end{bmatrix} = \begin{bmatrix} X & X & X & X & X & X \\ X & X & X & X & X & X \\ X & X & X & X & X & X \\ X & X & X & X & X & O \\ X & X & X & X & X & X \end{bmatrix} \begin{bmatrix} \text{DP1: AP Model} \\ \text{DP2: Location} \\ \text{DP3: Channel} \\ \text{DP4: Tx Power} \\ \text{DP5: Network Mgmt} \\ \text{DP6: Wired Network} \end{bmatrix}$$

Figure 8: Generic Wi-Fi Design Matrix.

In Axiomatic Design terms, the goal is to decouple the design, by selecting particular DPs so as to eliminate the coupling term in the design matrix. In Wi-Fi, while the DP categories themselves are fixed as defined above, a set of *best practices* can be derived and implemented for each DP, so as to minimize or eliminate the coupling terms.

While many of the Wi-Fi best practices highlighted below are performed by Wi-Fi network designers and troubleshooters, these practices have often been discovered and codified over years of trial and error. Furthermore, many of these are not necessarily universally accepted or adhered to amongst experts in the field. The best practices presented below are based upon the author's experience in the field, combined with careful consideration of typical Wi-Fi problems from an Axiomatic Design perspective.

It has been suggested that best practices can be treated as an explicit set of design constraints, since "FR creep" needs to be accommodated in a system design over the expected lifetime of the system. This was demonstrated in prior work on applying Axiomatic Design to naval ship design. [14] In the naval ship design case, the constraints specified are primarily focused on providing quantifiable margins in the design (i.e. that the DPs capabilities exceeded the FRs as specified), so as to accommodate expected changes to operational expectations (FRs) over the life of the naval vessel (DPs). While Wi-Fi design is clearly subject to the same type of margin constraints, the best practices themselves are not necessarily quantifiable, and may be impractical to follow depending upon other constraints in the system. As an example, it is a best practice to not deploy APs down a corridor, but rather in units on alternating sides of the central hallway. This may not be doable in the presence of a design constraint for an existing structure where cabling can only be cost-effectively run in the hallways and not into units. A good Wi-Fi design (i.e. one that satisfies the FRs and Cs) can still probably be generated, though it is likely to be more difficult and the resultant margins on the FRs will be lower. Accordingly, best practices are treated as guidelines, intended to advise the lower-level design choices as to how the DPs are implemented, such that the coupling terms in the design matrix are minimized or eliminated.

In order to decouple the design, a further understanding of the coupling terms in the design matrix is required, along with derivation of best practices to eliminate or minimize those coupling terms.

DP1: AP Model / Antenna

The AP model and antenna type (DP1) is chosen primarily for the usage type (FR1). That said, virtually by definition the particular model of AP and its integrated or external antennas will directly influence the requirements for coverage (FR2), capacity (FR3), control (FR4) and integration (FR5), and thus drive the particulars for their corresponding DPs, as shown in Figure 9.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3: Capacity} \\ \text{FR4: Control} \\ \text{FR5: Integration} \end{bmatrix} = \begin{bmatrix} X \\ X \\ X \\ X \\ X \end{bmatrix} [\text{DP1: AP Model}]$$

Figure 9: Design Matrix, portion for DP1.

DP2: AP Locations

The locations of the access points (DP2) are primarily chosen to satisfy the requirements for coverage (FR2). Depending on the AP model, multiple APs co-located in the same room are often required to provide sufficient capacity (FR3). The number and position of the APs also clearly impacts control (FR4) and integration (FR5). However, APs that are placed within direct line of sight and/or too close to each other will create co-channel interference, which will degrade network performance and thus the capability of the network to meet usage and application needs (FR1), resulting in a coupled design as shown in Figure 10.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3: Capacity} \\ \text{FR4: Control} \\ \text{FR5: Integration} \end{bmatrix} = \begin{bmatrix} X \\ X \\ X \\ X \\ X \end{bmatrix} [\text{DP2: Location}]$$

Figure 10: Design Matrix, portion for DP2.

Best practices to decouple the design: With respect to the client devices being serviced, each AP should be located as close as possible with the minimum number of obstructions. With respect to other APs, we want to discourage overlapping communications, so the APs should be placed as far apart as possible with as many intermediate obstructions as possible. In practical terms, this means alternating the position of APs both horizontally and vertically. APs should be placed in rooms on alternating sides of the hallway and staggered from floor to floor. By positioning APs so that they have minimal AP-to-AP interference, the impact of the AP position on ultimate performance is minimized. [9,10,11]

This is illustrated in the following example of a Wi-Fi design for a multi-level hotel. In Figure 11, the APs are located in the hallways, and while staggered from floor to floor, they still result in coverage problems in particular guest rooms, as well as high co-channel interference between APs, despite having proper channel (DP3) and transmission power (DP4) settings.

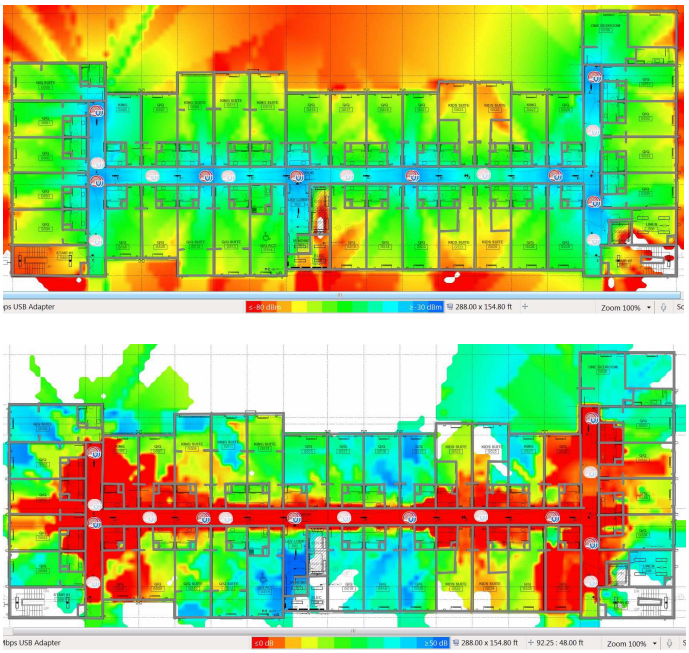

Figure 11: Bad design locations: Predictive coverage model (top) and signal-to-interference ratio (bottom) for hotel with APs in hallways (faded APs are on neighboring floor).

By contrast, Figure 12 shows the result of placing the same number of APs in alternating rooms on each side of the hallway. Coverage is dramatically improved, and co-channel interference between APs, while not eliminated because of the limited channel choices on the 2.4 GHz band, is drastically reduced.

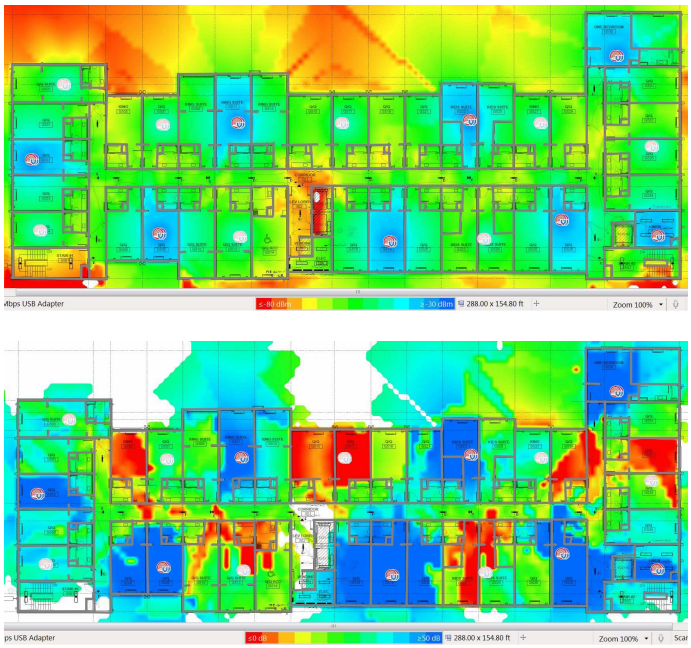

Figure 12: Good design locations Predictive coverage model (top) and signal-to-interference ratio (bottom) for hotel with APs in rooms (faded APs are on neighboring floor).

DP3: Channel

The channel width and center frequency of the access points (DP3) are primarily chosen to satisfy the requirements for capacity (FR3), while influencing the requirements for control (FR4) and integration (FR5). Given the restrictions on the number of independent channels in each band and how that decreases as the channel width increases, poor channelization will create AP-to-AP interference and thus degrade both usage (FR1) and coverage (FR2). This results in a coupled design as shown in Figure 13.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3: Capacity} \\ \text{FR4: Control} \\ \text{FR5: Integration} \end{bmatrix} = \begin{bmatrix} X \\ X \\ X \\ X \\ X \end{bmatrix} [\text{DP3: Channel}]$$

Figure 13: Design Matrix, portion for DP3.

Best practices to decouple the design: On the 2.4 GHz band, only using 20 MHz channel sizes should be used, and channels should be deployed across APs with an alternating static 1, 6, 11 scheme, both horizontally and vertically. For the 5 GHz band, the use of 160 MHz channel widths should be avoided, and an alternating static channel scheme should be deployed across APs, based upon Figure 7 for the given channel width. To keep capacity consistent throughout all APs on the network, all APs throughout the network should use the same channel width.

Many enterprise vendors also provide a feature called *band steering*, which encourages dual band client devices to connect

over the 5 GHz band to obtain higher speeds and lower interference than the 2.4 GHz band. Band steering should generally be enabled to provide such functionality.

It should also be noted that many vendors provide a feature either within the AP itself or on the controller called *auto channel* to automatically select the channels used by each AP. Each vendor implements this feature differently, though it is generally intended to make Wi-Fi deployments easier, and to react to changes in interference from the external environment by not requiring a static channel plan. The problem with this approach is that these algorithms, while great for marketing, tend not to work in actual practice, and in fact can make the network performance worse. Most vendors use methods whereby each AP performs a periodic brief scan of all of the channels to see what else is operating on each channel, and then selects the least noisy channel. This approach is fundamentally flawed: not only does the AP not get a true understanding of the channel usage over time, but it can be deceiving, especially on the 2.4 GHz band, where a channel can be seen as "clear", even when there is a lot of traffic on overlapping channels. Such methods, depending on the specific implementation and search algorithms, also tend to be convergent, meaning that neighboring APs will settle on the same or overlapping channels, thereby increasing co-channel interference between APs. Some vendors have proposed more sophisticated methods [12], though these alternatives tend to be divergent, i.e. change channels very frequently and not settle down. It is therefore considered best practice to turn off auto channel and use static channel plans for both the 2.4 GHz and 5 GHz bands.

DP4: Transmit Power

The transmit power (DP4) primarily addresses the size of coverage (FR2), while influencing the requirements for control (FR4) and integration (FR5). However, as discussed above, transmit power can be mismatched between the client devices and the access point, causing degradation in usage (FR1). Furthermore, as the cell size gets larger, the ability to re-use channels decreases and the likelihood of AP-to-AP interference increases, degrading capacity (FR3). This results in a coupled design as shown in Figure 14.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3: Capacity} \\ \text{FR4: Control} \\ \text{FR5: Integration} \end{bmatrix} = \begin{bmatrix} X \\ X \\ X \\ X \\ X \end{bmatrix} [\text{DP4: Tx Power}]$$

Figure 14: Design Matrix, portion for DP4.

Best practices to decouple the design: Since there are more DPs than FRs, the transmit power is the most natural choice to designate as the redundant DP. Accordingly, power settings for each band should be set uniformly across all APs on the network, and only tweaked marginally on individual APs at the end of an installation to boost coverage in particular areas. The specific transmit power settings are directly a function of the limitation of the client devices on the network and the antenna gain (i.e. AP model), so these values should be selected once the specific AP make and model are known. Furthermore, there should be a 4-6 dB transmit power offset between the 2.4 GHz band and the 5 GHz band to account for differences in signal propagation across the two bands.

It should also be noted that many vendors provide a feature either within the AP itself or on the AP controller called *auto power* to dynamically set the transmit power levels used by each AP. Each vendor implements this feature differently, though it is generally intended to make Wi-Fi deployments easier, and to react to changes in interference from the external environment by adjusting the power to minimize co-channel interference. As with auto channel, auto power tends not to work very well in practice, and in fact can also make the network performance worse. If the transmit power of an AP changes, it will change its effective coverage area. If transmit power is decreased, coverage (FR2) may be lost in particular areas, while increasing transmit power will increase coverage overlap, thus making co-channel interference between APs more likely, degrading usage (FR1) and capacity (FR3). Hence, it is considered best practice to not use auto power, but to set each AP with uniform power settings and only make minor adjustments as needed.

DP5: Network Management System

As the DP for addressing system control (FR4), the network management system (DP5) should be designed to not have influence over the primary functional requirements of usage (FR1), coverage (FR2), and capacity (FR3). [4] How well a particular network management system adheres to this depends on the particular architecture (i.e. central, distributed, or split), and whether features like auto channel and auto power are enabled or not. See Figure 15.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3: Capacity} \\ \text{FR4: Control} \\ \text{FR5: Integration} \end{bmatrix} = \begin{bmatrix} X \\ X \\ X \\ X \\ X \end{bmatrix} [\text{DP5: Network Mgmt}]$$

Figure 15: Design Matrix, portion for DP5.

Best practices to decouple the design: Generally, a distributed or split architecture for the network management system should be implemented, to prevent the AP controller from becoming a bottleneck in the system. Furthermore, the features and configuration settings selected in the AP configuration should be consistent with the best practices for the primary DPs. This will mean that some particular features, such as band steering, should always be enabled, whereas other features, such as auto channel and auto power, should always be disabled.

DP6: Wired Network

As the DP for addressing system integration (FR5), the wired network (DP6) must be designed so as to not have

influence over the primary functional requirements of usage (FR1), coverage (FR2), and capacity (FR3), and control (FR4). [4] How well a particular wired network adheres to this depends on the particular architecture and the components used on the wired side of the network. The primary drivers here will be the capacity of the Ethernet wiring and switches. See Figure 16.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3: Capacity} \\ \text{FR4: Control} \\ \text{FR5: Integration} \end{bmatrix} = \begin{bmatrix} X \\ X \\ X \\ X \\ X \end{bmatrix} [\text{DP6: Wired Network}]$$

Figure 16: Design matrix, portion for DP6.

Best practices to decouple the design: The design impact can be decoupled here by following Wi-Fi Design Best Practices and Design for System Architecture. The bottleneck of the network, in terms of sustained throughput, should be the bandwidth coming into the property by the carrier. Thus, it is important that the switch and cabling architecture have more bandwidth capacity than the bandwidth into the property or what the access points can provide. This may mean using higher capacity / higher cost components, utilizing link aggregation to increase the capacity of the links between network switches, utilizing spanning tree to purposely wire in loops into the network for redundancy, etc. It also means selecting wired components that have specific MAC layer 2 features to support the primary FRs, such as VLAN and access control list (ACL) capability.

Using the best practices identified above, the design matrix now becomes decoupled. The DPs are re-ordered slightly to move the redundant DP of transmit power (DP4) immediately after the selection of the AP model and antenna (DP1). This is shown in Figure 17.

$$\begin{bmatrix} \text{FR1: Usage} \\ \text{FR2: Coverage} \\ \text{FR3:Capacity} \\ \text{FR4: Control} \\ \text{FR5:Integration} \end{bmatrix} = \begin{bmatrix} X & X & 0 & 0 & 0 & 0 \\ X & X & X & 0 & 0 & 0 \\ X & X & X & X & 0 & 0 \\ X & X & X & X & X & 0 \\ X & X & X & X & X & X \end{bmatrix} \begin{bmatrix} \text{DP1: AP Model} \\ \text{DP4: Tx Power} \\ \text{DP2: Location} \\ \text{DP3: Channel} \\ \text{DP5: Network Mgmt} \\ \text{DP6: Wired Network} \end{bmatrix}$$

Figure 17: Design matrix, decoupled with best practices

TROUBLESHOOTING EXAMPLE

Figure 18 shows the floor plan for a multi-level hotel that the author was called in to troubleshoot. The hotel consisted of six floors, made up of drywall (walls) and gypsum (ceiling / floor), both of which are relatively transparent materials for RF. Five single band 2.4 GHz 802.11n access points were mounted in the hallways on each floor, in the same positions on each floor. The APs were all set to the maximum power output (29 dBm / 800 mW), and auto channel was enabled. Hotel guests were complaining about intermittent speeds and dropped connections.

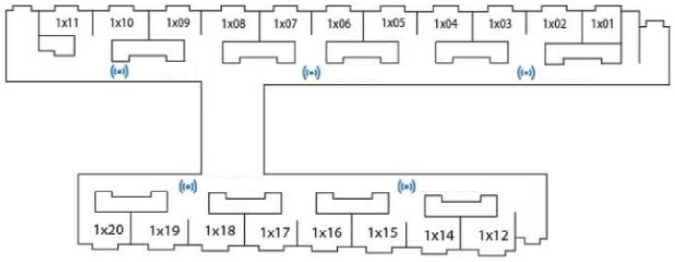

Figure 18: Hotel with poorly implemented DPs.

When initially evaluating the network, a scan from one of the APs on the first floor was done, and all of the other APs in the structure were observed, with power levels ranging from -40 dBm to -95 dBm. As a general guide, neighboring APs on non-overlapping channels should not exceed -60 dBm and neighboring APs on the same or overlapping channels should not exceed -82 dBm.

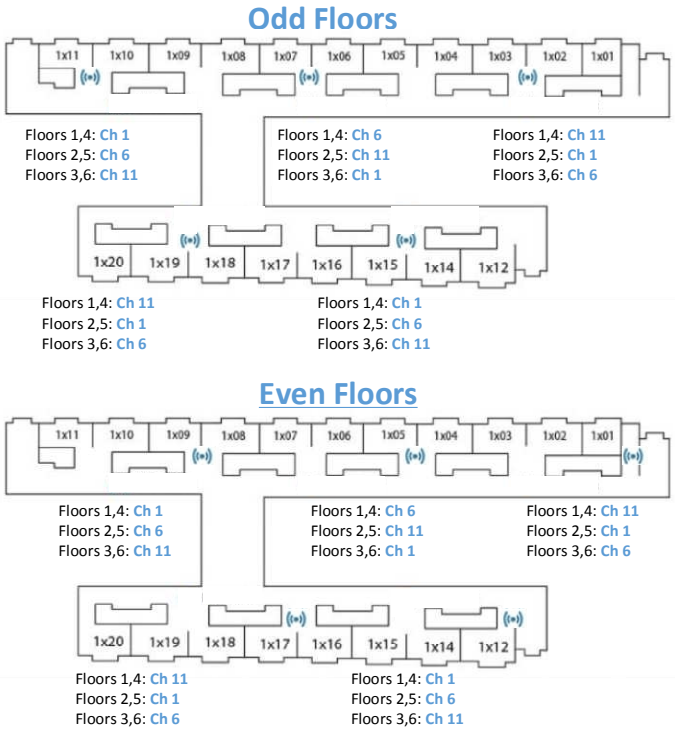

Figure 19: Hotel with DPs corrected for best practices, subject to constraints.

To remediate the issues, best practices were imposed to decouple the parameters. The hotel had drop ceilings in the hallways but hard ceilings in the guest rooms, so the APs could not be placed in the rooms themselves. Specifically, the following actions were taken, as shown in Figure 19.

- The APs were physically moved into alcoves in front of each pair of rooms, and their positions were staggered from floor to floor.
- A fixed 1,6,11 channel scheme was imposed.

- Transmit power levels were reduced to 20 dBm. (For two APs, power was increased to 23 dBm to address coverage issues in corner rooms).

The implementation of Wi-Fi best practices resolved the customer issues, as the co-channel interference between all of the APs within the hotel was dramatically reduced.

CONCLUSION

The design of a high performance Wi-Fi network is a complex engineering task subject to ever-increasing demands on its requirements and constraints. As such, design tools are needed to identify and validate best practices in Wi-Fi design, and to troubleshoot problems in existing deployments, and thus maximize the performance of the network.

This paper has applied Axiomatic Design to the problem of Wi-Fi network design and troubleshooting. Three primary independent FRs were identified: usage type, signal coverage, and capacity, with four corresponding DPs: AP model / antenna, location, channel, and transmission power. Additionally, two FRs were incorporated for command & control and support & integration of the system, with the corresponding DPs for network management and wired network infrastructure.

As Axiomatic Design illustrates, Wi-Fi systems are inherently coupled, due to the excessive number of DPs and the intrinsic constraints that are imposed between the FRs. Nonetheless, a set of Wi-Fi best practices can be defined in the proper definition of the DPs, so as to control the cross-coupling terms and decouple the design matrix. Applying these best practices and configuring the DPs in the proper order leads to more robust and properly functioning Wi-Fi systems.

ACKNOWLEDGMENTS

I dedicate this paper to my former teacher, advisor, and mentor, Professor Nam P. Suh. Professor Suh taught me a unique method of thinking and approaching engineering problems and challenges that has repeatedly proved to be invaluable across multiple engineering disciplines throughout my career.

REFERENCES

[1] Suh, N. P. *The Principles of Design*. Oxford University Press, New York. ISBN 019-504345-6. Copyright 1990.

[2] Suh, N. P. *Axiomatic Design: Advances and Applications.* Oxford University Press, New York. ISBN 019-513466-4. Copyright 2001.

[3] Tate, D. "A Roadmap for Decomposition: Activities, Theories, and Tools for System Design." Ph.D. Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA. February 1999.

[4] Hintersteiner, Jason D. "A Fractal Representation for Systems." *Proceedings of the 1999 International CIRP Design Seminar*, Enschede, the Netherlands. March 24-26, 1999.

[5] Hintersteiner, J. D. and Nain, A. "Integrating Software into Systems: An Axiomatic Design Approach." *Proceedings of the 3rd International Conference on Engineering Design and Automation*, Vancouver, B. C. Canada. August 1-4, 1999.

[6] Hintersteiner, J. D. and Tate, D. "Command and Control in Axiomatic Design Theory: Its Role and Placement in the System Architecture." *Proceedings of the 2nd International Conference on Engineering Design and Automation*, Maui, HI. August 9 – 12, 1998.

[7] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update 2014-2019 White Paper". http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html. February 3, 2015.

[8] "Gartner Says the Internet of Things Installed Base Will Grow to 26 Billion Units By 2020." http://www.gartner.com/newsroom/id/2636073. December 12, 2013.

[9] Coleman, D. and Westcott, D. *CWNA Certified Wireless Network Administrator Official Study Guide: Exam CWNA-106.* 4th edition. John Wiley & Sons, Inc., Indianapolis, IN. ISBN 978-1-118-89370-8. Copyright 2014.

[10] Jackman, S., Swartz, M., et al. *CWDP Certified Wireless Design Professional Official Study Guide: Exam PW0-250.* . John Wiley & Sons, Inc., Indianapolis, IN. ISBN 978-0-470-76904-1. Copyright 2011.

[11] Hintersteiner, J. *EnGenius Certified Operator.* EnGenius Technologies, Inc. certification program course. http://tinyurl.com/p9kfenc Copyright 2014-2015.

[12] Ruckus Wireless, Inc. *ChannelFly: Predictive Capacity Management for Automatic RF Channel Selection.* http://c541678.r78.cf2.rackcdn.com/feature-sheets/fs-channelfly.pdf. Copyright 2012.

[13] Suh, N.P. *Complexity: Theory and Applications*, Oxford University Press, 2005, ISBN 0-19-517876-9. Copyright 2005.

[14] Whitcomb, C.A. and Szatkoski, J.J. "Concept Level Naval Surface Combatant Design in the Axiomatic Approach to Design Framework." *Proceedings on ICAD2000 – First International Conference on Axiomatic Design*, Cambridge Ma. June 21-23, 2000.