

we cannot yet implement and evaluate phase-switched flyways, we ignore switching overhead in this paper and instead focus on individual traffic matrices in isolation.

Discovering the steering coefficients: Steering coefficients allow directional antennas of the sender and transmitter to point at each other. Optimizing the search process to discover steer coefficients is currently an active area of research. However, the flyways scenario greatly simplifies this problem. We use the wired network to coordinate the steering. Rather than needing to optimize the combination of multiple reflections off walls and in-home objects, in the DC, we can use physically meaningful steering patterns that point in a particular direction. A movement of a few mm does not dramatically reduce gains. The stable DC environment allows us to use history and retrain infrequently. Any nodes not involved in flyways (there will be many of these) can opportunistically measure their directional gains with respect to ongoing transmissions, as well as update measurements for patterns between idle nodes that will not interfere with ongoing traffic.

B. 60 GHz SIMULATOR

We implemented an 802.11ad simulator in `ns-3`. To have confidence that our simulations are a good reflection of reality, we base wireless effects directly on the physical layer measurements we took in §3 and the WiGig/802.11ad PHY and MAC design [31]. Here we describe wireless aspects of our `ns-3` model. We extended `ns-3` with other support too, such as automatic generation of DC layouts and routing, but these components are straightforward and we omit them due to lack of space.

Directional antennas: We built table-driven models from the measured radiation patterns of the antennas in our lab (Figure 3). We interpolate between measurements when needed. As well as using measured patterns, rather than the manufacturer antenna specifications, we take care to simulate the full 360° radiation pattern, not just the primary lobe. We also added a simple isotropic antenna model, and the radiation pattern used for Geo-fencing [25].

IEEE 802.11ad PHY and MAC: We implemented in `ns-3` the physical and MAC layers defined in the draft 802.11ad standard. We limit ourselves to the faster OFDM PHY. We fix transmit power to 10 mW to match commercial devices.

Signal propagation: We model signal propagation using Friis’ law. Our measurements (§3.2) show that this is a good fit for line-of-sight environments. Still, we conservatively subtract an additional 3 dB from the signal power (but not from interference) to represent potential destructive multi-path interference received via side lobes.

Interference (SINR): To calculate the SINR needed for bit error rate estimation, `ns-3` uses the standard SINR modeling technique. It adds together the power from multiple interferers, combines it with noise, and compares it with signal strength. `ns-3` does not model symbol-level fading, i.e., it assumes that the received power (RSS) from each transmitter is consistent throughout its transmission. It does, however, compute different SINR levels for different parts of packets when interference stops or starts during reception.

Our measurements of the stability of real links (§3.3) show that we can use this SINR model and ignore fading at the sub-packet level. Prior work (DIRC [17]) has also found this simple SINR model to be appropriate with directional antennas, even when using the 802.11g OFDM rates in non-line-of-sight environments and omni-directional antennas at receivers. The model is much more fitting in our 60 GHz domain: both transmitter and receiver use directional antennas so that secondary rays (multi-path) have little impact (§3.2); and the channel is very stable due to little environ-

| Ideal Rate | Wireless TCP | Offload ACKs to Wired | No DCF |
|------------|--------------|-----------------------|-----------|
| 693 Mbps | 656 Mbps | 672 Mbps | 676 Mbps |
| 6.76 Gbps | 4.58 Gbps | 5.36 Gbps | 5.62 Gbps |

Table 2: Impact of sending TCP ACKs over wire

mental mobility (§3.3).

Bit error rate (BER): Estimating the bit error rate, and hence whether a transmitted packet is received correctly, forms a key function of any wireless model. The input to this calculation is the SINR and the 802.11ad wireless rate. To estimate BER, we use the 802.11ad standard as our guide. It defines the sensitivity for each rate and coding as the (SINR) power level down to which a device much successfully receive more than 99% of 4096-byte packets sent using that rate. This reception rate corresponds to a BER less than 3.07×10^{-7} , and thus we calibrate our error model for each rate by assuming its BER is 3×10^{-7} when its SINR is the sensitivity threshold. The sensitivities defined in the standard implicitly include the (≈ -81 dBm) thermal noise for a 2.16 GHz channel, and a 15 dB combined implementation loss. We compute BERs at other SINR values using textbook formulas [23] for BER as a function of SNR in Gaussian noise. To receive a packet, all bits must be correct.

Auto-rate algorithm: The 802.11ad standard does not mandate use of a specific auto-rate algorithm. We select rates based on received SINR. This is reasonable for our stable data center environment.

C. IMPROVING WIRELESS PERF.

Data center performance will improve the most when the flyways deliver the largest possible throughputs. A unique aspect of our flyways scenario is the hybrid wired and wireless nature of the network, and in this section we describe and evaluate two wireless optimizations that leverage the wired backbone in the DC to increase flyway TCP throughput by 25%

Wired offload of MAC-inefficient packets: TCP ACKs are far smaller than data packets, and make inefficient use of wireless links because payload transmission time is dwarfed by overheads such as preamble and SIFS. The hybrid wired-wireless design of our network lets us improve efficiency by sending ACK packets over the wire instead. To measure the improvement, we simulated a single TCP flow on a 20 m link and configured `ns-3` to send the TCP ACKs over the wired network. Table 2 shows the resulting TCP throughput. For fast links enabled by the narrow-beam antenna, the performance improves 17%. Note that the TCP ACK traffic will use some wired bandwidth, but this will be trivial compared to the increase in throughput.

Removing DCF: For the common case of one-way TCP flows in the data center [4], if we divert TCP ACKs over the wire as above then all traffic over a given wireless link will flow in only one direction. Furthermore, our system design (§5) is based on independent flyways that do not interfere with one another. Thus, there are no collisions in our wireless network, and we can eliminate the DCF backoff mechanism. This change improves the TCP throughput by an additional 5%, as seen in the third column of Table 2.

Occasionally there may be bidirectional data flows over the flyway. Even in this case, we can remove the cost of DCF. Since only the two communicating endpoints can interfere with each other, we can easily schedule transmissions on the link by passing a token between the endpoints. This naturally fits into the 802.11 link layer protocol because after transmitting a packet batch, the sender waits for a link layer Block-ACK. We can exploit this scheduled hand-off to let the receiver take the token and send its own batch of traffic.