# Beamforming on Mobile Devices: A First Study

Hang Yu, Lin Zhong, Ashutosh Sabharwal, and David Kao
Department of Electrical & Computer Engineering, Rice University, Houston, Texas, USA
{hang.yu, lzhong, ashu, davidkao}@rice.edu

## ABSTRACT
In this work, we report the first study of an important realization of directional communication, *beamforming*, on mobile devices. We first demonstrate that beamforming is already feasible on mobile devices in terms of form factor, device mobility and power efficiency. Surprisingly, we show that by making an increasingly profitable tradeoff between transmit and circuit power, beamforming with state-of-the-art integrated CMOS implementations can be more power-efficient than its single antenna counterpart. We then investigate the optimal way of using beamforming in terms of device power efficiency, by allowing a dynamic number of active antennas. We propose a simple yet effective solution, *BeamAdapt*, which allows each mobile client in a network to individually identify the optimal number of active antennas with guaranteed convergence and close-to-optimal performance. We finally report a WARP-based prototype of BeamAdapt and experimentally demonstrate its effectiveness in realistic environments, and then complement the prototype-based experiments with Qualnet-based simulation of a large-scale network. Our results show that BeamAdapt with four antennas can reduce the power consumption of mobile clients by more than half compared to a single antenna, while maintaining a required network throughput.

## Categories and Subject Descriptors
C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design - Wireless Communication

## General Terms
Algorithm, Design, Experimentation, Measurement

## Keywords
Beamforming, Mobile devices, BeamAdapt, Power efficiency

## 1. INTRODUCTION
Current and emerging wireless networks all assume that their mobile accessing clients are omni directional in uplink transmission, radiating power equally toward all directions. As the number of mobile devices explodes, such omni directionality not only limits the device power efficiency due to the waste of radiation power, but also becomes a critical barrier to the network capacity due to interference between peer clients. In this work, we study a client-based approach toward addressing the omni directionality: *using beamforming on mobile devices for directional transmission*. By focusing the transmit power toward the proper direction, beamforming can not only improve the signal to noise radio (SNR) at the intended receiver but also reduce the interference engendered to the others.

While beamforming has been well studied and already deployed on base stations, access points, and vehicular platforms, it is never examined on mobile devices due to three physical challenges of the mobile device: small size, high mobility, and limited power. Naturally, the first question one may ask is: *is beamforming feasible on mobile devices?* We answer this question by examining the three challenges (Section 3). First, we show that single-user beamforming with two to four antennas can fit into mobile devices with a linear or circular array. Second, we experimentally demonstrate that the beamforming gain remains high even when the device can not only move but also rotate at a high speed. Finally, using data from research prototypes and emerging products, we show that beamforming can be even more power-efficient than its single antenna counterpart, by making an increasingly profitable tradeoff between transmit and circuit power. More importantly, the power tradeoff made by beamforming leads to an optimal number of active antennas, or an optimal *beamforming size*, which minimizes the device overall power. We reveal that the optimal beamforming size is dependent on the channel condition and required link capacity, which strongly suggests an adaptive use of beamforming that optimizes the beamforming size and turns off idle antennas for power efficiency.

Such adaptive beamforming is straightforward to realize for a single link because the most power efficient beamforming size can be analytically calculated. However, with multiple interfering clients, identifying the optimal beamforming size for each client is challenging, due to not only the absence of an analytical solution, but also the requirement of client cooperation to enumerate all possible beamforming size combinations. Therefore, the second question one may ask is: *can each mobile client in a large-scale network individually identify its optimal beamforming size that collectively approaches the optimal tradeoff with minimum network power consumption?* We answer this question by proposing BeamAdapt, a distributed solution with which each client optimizes its beamforming size without coordination with others (Section 4). The key idea of BeamAdapt is to let each client iteratively adjust its beamforming size solely based on the SINR at its own receiver. Regardless of its simplicity, BeamAdapt has guaranteed convergence and closely approaches its performance bound, as shown by our empirical results.

We evaluate BeamAdapt first through a prototype-based experiment of a two-link network (Section 5) and then a Qualnet-based simulation of a large-scale network (Section 6). Our experimental results show that in various indoor and outdoor environments, BeamAdapt is able to tightly maintain the required SINR even with client mobility, and meanwhile consumes only 5% higher power compared to a genie-aided solution as the performance bound in reality. For our Qualnet-based simulation, we realize

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
*MobiCom'11*, September 19–23, 2011, Las Vegas, Nevada, USA.
Copyright 2011 ACM 978-1-4503-0492-4/11/09…$10.00.


Samsung Ex. 1024, Page 1 of 12

BeamAdapt in the context of modern cellular systems. We show that by leveraging uplink power control, one can easily realize BeamAdapt on mobile clients with trivial protocol modification. We show that within a large-scale cellular network, BeamAdapt with four antennas can reduce power consumption of the client wireless transceiver by 54%, while achieving similar network throughput.

In summary, we make the following technical contributions toward beamforming on mobile devices:

- We report the first feasibility study of beamforming on mobile devices in terms of form factor, device mobility and power efficiency. Our examination shows that beamforming is not only feasible but also profitable to mobile devices.
- We provide a concise solution, BeamAdapt, which allows each client in the network to rapidly identify the optimal beamforming size to achieve the required capacity with close-to-maximal power efficiency. The simplicity of BeamAdapt allows its immediate realization on current mobile devices.
- We report a prototype of BeamAdapt based on the WARP platform and a system design for realizing BeamAdapt on the clients of cellular networks. Our prototype-based and Qualnet-based evaluations collectively demonstrate the feasibility and power efficiency benefit of BeamAdapt.

BeamAdapt can be extended in two orthogonal ways (Section 7). First, while we propose BeamAdapt for transmit beamforming in this work, receive beamforming on mobile clients can similarly adopt BeamAdapt for power efficiency, with an even simpler formulation. Second, BeamAdapt leverages the beamforming gain to achieve client power efficiency given the capacity requirement. The beamforming gain can be alternatively used to improve the capacity given the client power constraint, indicating a dual formulation of BeamAdapt.

Although there has been recent research focus to enable directionality on mobile devices using passive directional antennas, e.g., [1, 2], our work is the first to study the possibility of power-efficient real-time beamforming on mobile devices in a large scale network. Beamforming, as a more flexible and more beneficial realization of directional communication, can fundamentally removes the limitation of client omni directionality on the network capacity and client power efficiency. We hope our initial study of beamforming on mobile devices can motivate and invite more serious research efforts to fully realize its potential in many other directions.

# 2. BEAMFORMING PRIMER

Unlike omni directional transmission with a single antenna, beamforming uses a group of antennas to increase the SNR of the received signal. Each antenna includes a passive antenna and a devoted RF chain that bridges the baseband signal and RF signal. Beamforming operates by adaptively assigning proper weights to the baseband signal and then transmitting the weighted signals through multiple antennas. It can be mathematically presented as:

$$\boldsymbol{x}(t) = \boldsymbol{w} \cdot s(t), \boldsymbol{x}(t) = \left(x_1(t), \cdots, x_N(t)\right), \boldsymbol{w} = (w_1, \cdots, w_N),$$

where the baseband signal, weight vector and output signal vector are denoted as $s(t)$, $\boldsymbol{w}$ and $\boldsymbol{x}(t)$, respectively.

The beamforming gain, $G$, is defined as the ratio of the received signal or interference power with beamforming, to that with a single antenna. Beamforming increases SNR and reduces interference by realizing a higher beamforming gain at the intended receiver than that at other receivers. Apparently, the number of active antennas, or the *beamforming size*, $N$, has a significant impact on the beamforming gain. With appropriate antenna spacing, beamforming with $N$ antennas can achieve a peak gain up to $N$, or $10\log(N)$ in dB [3], when signals from all transmit antennas add coherently at the intended receiver.

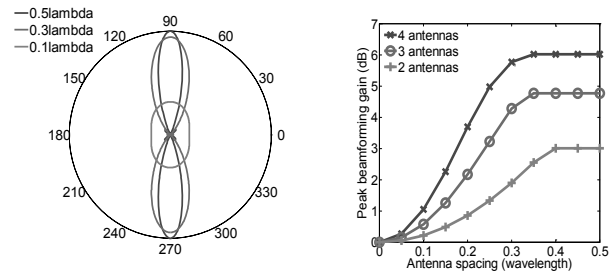


**Figure 1: (Left) Beamforming pattern of a linear array with four antennas, under different antenna spacing; (Right) Peak beamforming gain with different antenna spacing, for beamforming size from two to four. Both are for single-user beamforming.**

## 2.1 Antenna Spacing

Antenna spacing is critical for form-factor constrained mobile devices, and it directly affects the beamforming gain. In this work we only consider *single-user beamforming*, where the antenna spacing is relatively small (<0.5 $\lambda$ where $\lambda$ is the wavelength of the carrier signal) so that one cannot expect a diversity gain from the multiple antennas. As a result, the transmitter optimizes its weight vector to maximize the SNR for a single receiver. Figure 1 (Left) shows the beamforming pattern of a four-antenna linear array with different antenna spacing [3]. Apparently, when the antenna spacing decreases, the beamforming pattern becomes wider and the peak gain drops, which is also illustrated by Figure 1 (Right). We must note that Figure 1 is only about the transmit beamforming pattern, while the SNR at the receiver is additionally affected by the channel. The transmit beamforming pattern represents the beamforming gain relative to omni directional transmission under the same channel.

Figure 1 (Right) also shows that when the antenna spacing drops below certain threshold, the peak beamforming gain decreases, due to the power leakage toward a wider range of directions. The minimum antenna spacing for achieving the largest possible peak gain ($10\log(N)$ dB) depends on the number of antennas and is typically 0.3-0.4 $\lambda$. Again, we do not consider antenna spacing over 0.5 $\lambda$ since it is not only difficult to realize on form-factor constrained mobile devices but also undesirable due to its significant side lobes in single-user beamforming.

Multi-antenna techniques other than single-user beamforming usually have a more demanding requirement for antenna spacing. *Multi-user beamforming* [4] and *null beamforming* require the antenna spacing be above 0.5 $\lambda$ [5], in order to exploit additional degrees of freedom when choosing the weight vector. *Spatial multiplexing/diversity techniques*, a.k.a. MIMO techniques, typically need an antenna spacing of multiple wavelengths to operate with a satisfactory capacity improvement [5]. Apparently, they do not fit into smartphone-like mobile devices for the frequency bands in use today (2-5 GHz).



Samsung Ex. 1024, Page 2 of 12

## 2.2 Channel Estimation

To guarantee the signals from multiple transmit antennas add coherently at the receiver to achieve the maximal beamforming gain, beamforming requires channel knowledge at the transmitter. In single-user beamforming, the weight vector is assigned as $\boldsymbol{w} = \boldsymbol{h}^*$, where $\boldsymbol{h}$ is the channel vector in which each of its elements represents the corresponding coefficient of the channel between a transmit antenna and the receiver. The channel vector $\boldsymbol{h}$ is often generally denoted as *Channel State Information* (CSI).

For transmit beamforming, CSI can be obtained through either closed-loop or open-loop estimation. For closed-loop CSI estimation, the transmitter needs to send training symbols to the receiver, and then the receiver leverages the training symbols to calculate the channel coefficients and sends the CSI back to the transmitter. For open-loop CSI estimation, the transmitter estimates the reverse channel when receiving and assumes it for the channel of transmitting. Apparently, open-loop CSI estimation requires channel reciprocity to be effective.

## 2.3 Power Characteristic

For single-user beamforming, given $\boldsymbol{h}$, the weight vector $\boldsymbol{w}$ is also given without the need of any additional computation. This is different to other MIMO techniques which often need considerable signal encoding and processing even at the transmitter. As a result, single-user beamforming incurs little power overhead to the baseband processing and we next focus on its RF power characteristic. Figure 2 illustrates the major RF hardware components of a beamforming transmitter. The transmitter consists of multiple RF chains, each of which is connected to a passive antenna. When we say an antenna is active, we mean that the RF chain connected to the antenna is powered on. When an antenna is not in use, the corresponding RF chain can be powered off to conserve power.

Accurately modeling the power characteristic of wireless transceivers is known to be challenging, especially with various transceiver realizations. However, since all the beamforming transmitters have similar components (Figure 2), and our power saving solution is as simple as turning off a RF chain, very accurate power modeling is unnecessary. Therefore, in this work we follow the widely-accepted power model proposed by [6] with improved modeling for the power amplifier. As shown by Figure 2, the transmitter power consumption includes that of the circuitry shared by all active RF chains, i.e., the *frequency synthesizer*, denoted as $P_{Shared}$, and that of each active RF chain. The power contributed by each active RF chain can be further broken down to that by the *power amplifier*, and that by the rest of the chain, denoted as $P_{Circuit}$. We assume identical power amplifiers for all the RF chains and combine their power consumption, jointly denoted as $P_{PA}$ with the output power from the transmit antenna included. Clearly, $P_{PA}$ is dependent on the total transmit power, $P_{TX}$, while $P_{Circuit}$ is constant irrespective of $P_{TX}$.

We model $P_{PA}$ as $P_{PA}=P_{TX}/\eta$, where $\eta$ is the efficiency of the power amplifier. The efficiency $\eta$ is usually dynamic depending on the transmit power, and here we approximate $\eta$ as a linear function of $P_{TX}$ [7] but note that the power amplifier itself is not necessarily linear. As a result, the total power of the beamforming transmitter, $P$, can be fairly accurately modeled as

$$P = \frac{P_{TX}}{\eta(P_{TX})} + NP_{Circuit} + P_{Shared}. \quad (1)$$

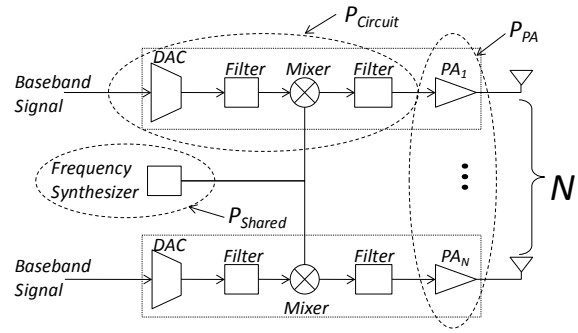


**Figure 2: RF components of a beamforming transmitter.**

In the rest of the paper, we adopt parameters as follows: $\eta_{min}$=0.3, $\eta_{max}$=0.5, $P_{Circuit}$=48.2 mW, $P_{Shared}$=50 mW. They are chosen based on [7, 8] as well as all recent CMOS wireless transceiver designs we have collected (see Section 3.3). Those parameters are on par with state-of-the-art transceiver designs in 2-5 GHz band [6].

# 3. FEASIBILITY STUDY

The first reaction one has toward beamforming on mobile device is likely to be: *is it feasible at all* (possibly thinking of the bulky, power-hungry Phocus Array system [9])? In this section, we examine three key physical challenges to put beamforming on mobile devices: form factor, device mobility and power efficiency. Our key conclusion after a careful examination is: beamforming not only is feasible for mobile devices with a reasonable size, but also can improve their power efficiency if used properly.

## 3.1 Form Factor

With the advancement of semiconductor technologies, multiple RF chains are already being integrated into a single wireless transceiver chip, e.g., [10]. Therefore, the form factor challenge introduced by beamforming only stems from its antenna spacing requirement. As discussed in Section 2.1, beamforming typically requires the antenna spacing to be higher than 0.3-0.4 $\lambda$, or 4.5-6 cm at 2 GHz. There is no obstacle for medium-size mobile devices such as tablets and NetBooks to embrace four antennas, in either a linear array or a circular array. Small-size mobile devices such as smartphones can accommodate two antennas in a linear array or four in a circular array.

It is also worth noting that multi-antenna solutions using passive directional antennas reported in [1] do not have much antenna spacing requirement, because only one directional antenna is active at any time. However, the solution requires all the directional antennas to be properly oriented, introducing a different and even larger form factor challenge.

## 3.2 Device Mobility

A mobile device can not only move but also rotate. Recent work has shown that beamforming with predefined beam patterns can cope with vehicular mobility very well, e.g., [11, 12]. However, real-time beamforming imposes a new challenge due to the requirement of accurate CSI, including only not the magnitude but also the phase of the channel coefficients which are largely affected by device rotation. Therefore, next we focus on evaluating the beamforming gain under device rotation, since rotation can



Samsung Ex. 1024, Page 3 of 12

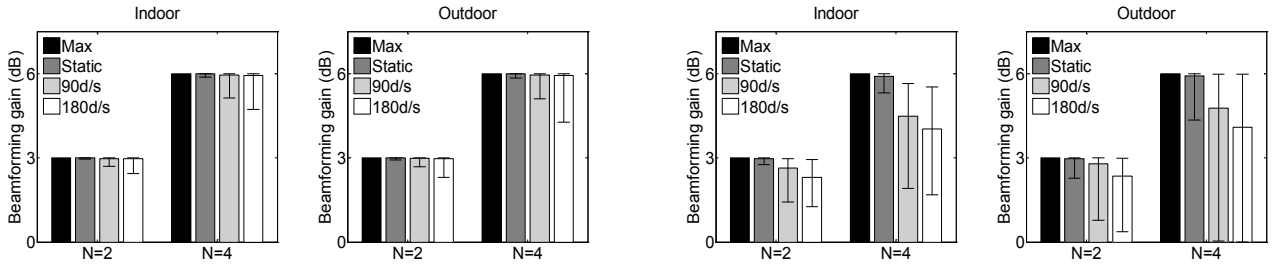


(a) CSI estimation per 10 ms (b) CSI estimation per 100 ms

**Figure 3: Beamforming gain under CSI estimation with various client rotation speeds. We show the results in different environments and with different CSI estimation frequencies.**

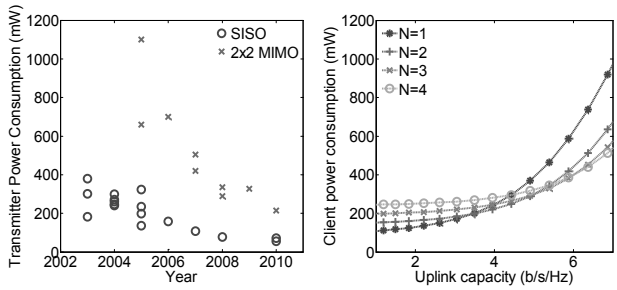


**Figure 4: (Left) Transmitter power trend from designs in ISSCC and JSSC; (Right) Client power consumption to deliver a range of link capacity.**

**Table 1: Simulation settings for the power tradeoff made by beamforming.**

| Parameters | Values |
|---|---|
| Distance | 0.5 km |
| Max beamforming size | 4 |
| Power decay factor | 4 |
| Receiver noise | -170 dBm/Hz |
| Channel bandwidth | 5 MHz |
| Carrier frequency | 2 GHz |

possibly introduce even faster channel variation than movement can.

We perform the experiments using the WARP software radios [13] and a concrete experimental setup is presented in the Appendix. The key question we aim to answer is: *what is the impact of device rotation on the CSI estimation and the corresponding beamforming gain?* To see this, Figure 3 shows the average beamforming gain under CSI estimation with different rotation speeds of the client node. In each sub-figure of Figure 3, four values of the beamforming gain for each beamforming size are shown: the upper bound given by perfect CSI (Max), the one given by estimated CSI with a stationary client (Static), the one given by estimated CSI with a rotating client at 180°/s (180d/s).

Clearly seen from Figure 3, when the CSI estimation interval is 10 ms, the CSI can be very accurate even with client rotation speed of 180°/s. As a result, the maximal beamforming gain, i.e., 3 dB and 6 dB with $N$=2 and $N$=4 respectively, can be closely achieved. When the interval is increased to 100 ms, the beamforming gain will be affected by client rotation. The rotation has higher impact for larger beamforming sizes due to a more focused beamforming pattern. Therefore, we conclude that even under high speed device rotation such as 180°/s, beamforming can still be effective with reasonable CSI estimation intervals, e.g., 10 ms. Finally, we observe that the performance of CSI estimation is more stable indoor, due to richer multipath effect to compensate faulty directions. This can be seen from the range of the beamforming gain in each sub-figure.

## 3.3 Power Efficiency

Compared to omni directional transmission with a single antenna, beamforming increases power consumption of the RF circuitry by simultaneously using multiple active RF chains. While the improvement of RF integrated circuits is slower than that of their digital counterparts, their power efficiency still improves significantly over years.

To illustrate this trend, we have examined the CMOS wireless transceiver realizations at 2-5 GHz, reported in ISSCC [14] and JSSC [15], the top conference and journal for semiconductor circuits, from 2003 to 2010. In Figure 4 (Left), we show the circuit power consumption, $P_{Circuit}+P_{Shared}$, of both single-antenna (SISO) and multi-antenna (MIMO) transceivers in their transmit mode. The figure clearly shows the continuous improvement in the power efficiency of both SISO and MIMO transceivers. As semiconductor process technologies continue to improve, $P_{Circuit}$ and $P_{Shared}$ will continue decreasing. Meanwhile, the power efficiency of power amplifiers is primarily limited by the transmit power, which will not change over time. As a result, $P_{PA}$ will increasingly dominate the total transmitter power consumption.

### *3.3.1 Power Tradeoff by Beamforming*

By focusing the transmit power toward the intended direction, beamforming can reduce the required transmit power, $P_{TX}$, and therefore the power consumption of the power amplifiers, $P_{PA}$. As a result, despite of higher circuit power consumption, beamforming is likely to improve the transceiver power efficiency. Clearly, beamforming makes a tradeoff between transmit and circuit power: with a beamforming size of $N$, the transmit power can be reduced to $1/N$ compared to a single antenna due to the beamforming gain. Note that beamforming is able to yield a total transmit power reduction instead of that of each antenna, i.e., the reduction is not because of the allocation of transmit power into multiple antennas.

We next briefly analyze this tradeoff between transmit and circuit power. For simplicity, we consider a single uplink channel from a



Samsung Ex. 1024, Page 4 of 12

mobile client to its infrastructure node and assume line-of-sight (LOS) propagation with the settings specified in Table 1. Figure 4 (Right) shows the client power consumption calculated by Equation (1) to deliver a range of link capacity for beamforming sizes from one to four. One can make two important conclusions from the figure. First, beamforming ($N$>1) is already more power-efficient than a single antenna ($N$=1) when delivering the uplink capacity of 3.2 b/s/Hz or higher. Second, the larger the required link capacity, the larger the most power-efficient beamforming size. Therefore, beamforming is increasingly desirable in delivering a higher capacity, which current wireless networks are trying to provide, in order to accommodate more mobile devices and throughput-intensive applications.

# 4. POWER-EFFICIENT BEAMFORMING ON MOBILE DEVICES

The above findings suggest an adaptive way to use beamforming on mobile devices: one shall adjust the beamforming size for the optimal tradeoff between transmit and circuit power, according to link capacity requirement. Next we show that to achieve the optimal tradeoff in a network is indeed non-trivial and, therefore, provide a solution, BeamAdapt.

## 4.1 Key Challenges

As shown in Figure 4 (Right), the optimal beamforming size varies according to the required link capacity. Given the power decay factor and distance, one can derive the required transmit power for omni directional transmission, $P_O$, to achieve certain link capacity. Using Equation (1), we can calculate the optimal beamforming size as

$$N_{opt} = \sqrt{P_O/C_1 P_{Circuit}} - C_2 P_O/C_1 \quad (2)$$

where $C_1$ and $C_2$ are constants determined by the power amplifier. Again, beamforming with more antennas is increasingly more efficient as $P_{Circuit}$ decreases according to the continual progress in semiconductor technologies.

While the optimal tradeoff given by $N_{opt}$ appears straightforward to identify with a single link ($P_O$ is uniquely decided by the required link capacity or SNR), it is challenging to determine in a network with multiple links. This is because $P_O$ is determined by SINR instead of SNR due to interference. Meanwhile, different beamforming size will generate different interference toward other receivers, implicitly affecting their own SINR. As a result, the optimal beamforming size can no longer be calculated by Equation (2).

Nonetheless, the tradeoff between transmit and circuit power is still valid and there exists a most power-efficient beamforming size for each client that collectively minimizes the aggregated client power consumption, or *network power consumption*. The immediate question we seek to answer is: *how could clients of a large network identify their most power-efficient beamforming sizes that collectively lead to the minimum network power consumption?*

## 4.2 Problem Formulation

We seek to minimize the aggregated power consumption by all the clients in a network, $P_{Network}$, with a constraint on the capacity, or equivalently the SINR of each link $i$, $SINR_i$. We separately constrain the SINR of each link since different links usually have different capacity requirements. In addition, the beamforming size, $N_i$, must be integers no greater than $N_{i,max}$, where $N_{i,max}$ is the number of antennas on client $i$.

**Algorithm 1: Identify the optimal beamforming size and transmit power for each client by BeamAdapt**

**Input**: SINR constraint $\rho$, max beamforming size $N_{max}$
**Output**: optimal beamforming size $N^{opt}$, transmit power $P_{TX}^{opt}$

$\left(P_{TX}^{(0)}, N^{(0)}\right) = (P_{TX}^{(0)}, 1)$, $k = 0$
obtain $SINR^{(0)}$
**while** $\left|SINR^{(k)} - \rho\right| \geq \varepsilon$
  $P_{min} = +\infty$
  **for** $N^{(k)} \leq N^{(k+1)} \leq N_{max}$
    compute $P_{TX}^{(k+1)}$
    $P = P_{TX}^{(k+1)}/\eta + N^{(k+1)} P_{Circuit} + P_{Shared}$
    **if** $P \leq P_{min}$
      $\left(P_{TX}^{opt}, N^{opt}\right) = \left(P_{TX}^{(k+1)}, N^{(k+1)}\right)$
      $P_{min} = P$
    **end**
  **end**
  obtain $SINR^{(k+1)}$
  $k = k + 1$
**end**
**return** $\left(P_{TX}^{opt}, N^{opt}\right)$

Therefore, we formulate the optimization problem as:

$$\text{minimize } P_{Network} = \sum_{i=1}^{M} P_i\left(P_{TX,i}, N_i\right)$$

$$\text{s.t. } SINR_i(\boldsymbol{P}_{TX}, \boldsymbol{N}) = \rho_i, 1 \leq N_i \leq N_{i,max}$$

where $P_i$ is the power consumption of client $i$ and

$$\boldsymbol{P}_{TX} = \left(P_{TX,1}, \cdots, P_{TX,M}\right), \boldsymbol{N} = (N_1, \cdots, N_M).$$

Solving this optimization problem is very challenging. First, each SINR constraint is a function of all $2M$ optimization variables. The SINR function is non-convex with respect to these variables, yielding the non-convexity of the problem. Second, there is no closed-form formulation of the beamforming gain to unintended receivers as a function of $N$. Its dependence on the receiver direction makes low-order approximation infeasible. Last, the integer constraint on $N_i$ renders a NP-hard mixed integer programming (MIP) problem [16]. While an exhaustive searching algorithm can ultimately offer the solution, the complexity can be as high as $O\left(\prod_{i=1}^{M}(N_{i,max})\right)$, which becomes prohibitive as $M$ grows. More importantly, such brute-force algorithm requires all the clients have knowledge of each others' actions in order to enumerate all beamforming size combinations, and cooperatively choose their beamforming sizes. Coordination between clients is known to be hard and overhead-intensive in wireless networks.

To tackle this, we introduce a distributed algorithm, BeamAdapt, with which each client simply performs individual optimization on the beamforming size without coordination.

## 4.3 Distributed Algorithm: BeamAdapt

First we decompose the problem into multiple, individual subproblems, i.e., the $i$th link's problem ($i$=1,2,⋯$M$) is formulated as

$$\min P_i \text{ s.t. } SINR_i = \rho_i.$$



Samsung Ex. 1024, Page 5 of 12

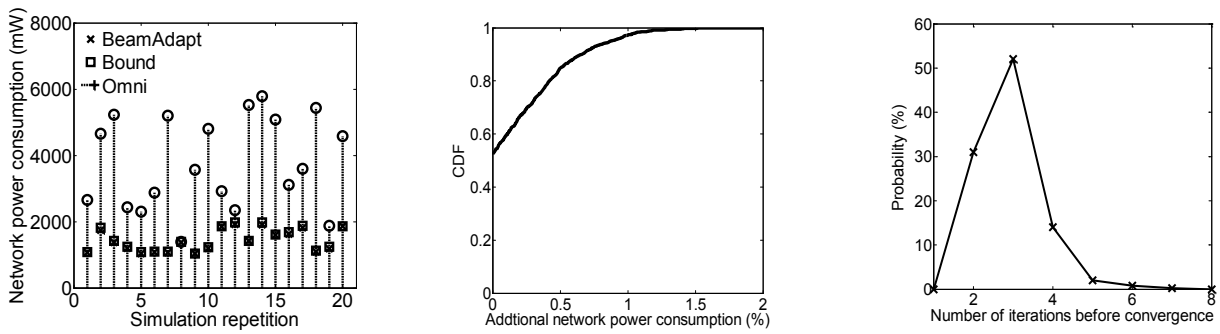


**(a) Network power consumption comparison** **(b) CDF of the additional power consumption by BeamAdapt** **(c) PDF of the convergence speed of BeamAdapt**

**Figure 5: Empirical results for the performance bound and convergence speed of BeamAdapt, in the seven-cell network.**

The optimal ($P_{TX,i}$, $N_i$) is determined iteratively. Let us temporally omit the subscript $\underline{i}$ below since all clients employ the same algorithm. We assume the transmit power and beamforming size are $P_{TX}^{(k)}$ and $N^{(k)}$ in the $k$th iteration, and the received SINR is $SINR^{(k)}$, then in the ($k$+1)th iteration, $P_{TX}^{(k+1)}$ and $N^{(k+1)}$ are identified by solving the following optimization problem:

$$\min_{P_{TX}^{(k+1)},N^{(k+1)}} P_{TX}^{(k+1)}/\eta + N^{(k+1)}P_{Circuit} + P_{Shared}$$

$$\text{s.t. } \frac{P_{TX}^{(k+1)}N^{(k+1)}}{P_{TX}^{(k)}N^{(k)}} = \frac{\rho}{SINR^{(k)}},\ N^{(k+1)} \geq N^{(k)}.$$

The initial beamforming size is set to one, i.e., $N^{(0)} = 1$, while $P_{TX}^{(0)}$ can be arbitrary.

The iteration stops when $\left|SINR^{(k)} - \rho\right| \leq \varepsilon$, where $\varepsilon$ can be set according to the accuracy requirement. In each iteration, $(P_{TX}^{(k)}, N^{(k)})$ can be obtained by searching among all feasible beamforming sizes, with the complexity of $O\big(\max(N_{i,max})\big)$. Algorithm 1 shows the pseudo-code of the BeamAdapt. We note that when $M$=1, the problem reduces to single-link optimization which offers the same solution as Equation (2) does.

## 4.4 Convergence of BeamAdapt

The iteration process of BeamAdapt is guaranteed to converge. Next we provide a brief yet sufficiently illustrative proof. The two key facts we leverage are: (*i*) whenever the beamforming sizes are fixed, the iteration of BeamAdapt is isomorphic to a distributed power control algorithm that ensures convergence; (*ii*) the change of the beamforming size $N$ of each client is monotonous. That is, the beamforming size can only increase during the iteration.

Therefore, we divide the iteration process into multiple stages, $k_l (1 \leq l \leq L)$, where during each stage $N$ is constant and only $P_{TX}$ changes. The current stage $k_l$ evolves into $k_{l+1}$ when $N$ changes for any one link. Based on the monotonicity of $N$ we have the following inequality

$$L \leq \prod_{i=1}^{M}(N_{i,max}) < +\infty,$$

which indicates a finite number ($L$) of stages.

During each stage, the beamforming size is fixed; therefore the original problem turns in to

$$\min\ P_{Network} = \sum P_i(P_{TX,i}),\ \text{s.t. } SINR_i(\boldsymbol{P}_{TX}) = \rho_i.$$

This problem is isomorphic to the well-studied network power control problem where a distributed algorithm ensures convergence [17]. As a result, during each stage ($k_l$) the power control component either converges, or it moves onto a new stage. Since the number of potential stages $L$ is finite, the overall algorithm is guaranteed to converge.

## 4.5 Performance Bound of BeamAdapt

We next investigate the steady-state performance of BeamAdapt. It is possible that BeamAdapt converges to a sub-optimal solution. Unfortunately, the performance bound of BeamAdapt is not analytically obtainable, again due to the non-convexity of the optimization problem and the integer constraints on the beamforming size. Therefore, we have to rely on empirical methods to study the performance bound. We employ a seven-cell network that includes one hexagon cell and its six immediate neighbors of identical size. Each cell has an infrastructure node in the center that serves one client inside the cell at a time. Such seven-cell network configuration is often used as the first-order approximation of large-scale infrastructure networks. Other settings are similarly adopted from Table 1. To eliminate the dependency of BeamAdapt on the client location, we repeat the simulation extensively with random client locations. Therefore, we are in fact averaging the performance of BeamAdapt with dynamic network configurations.

Figure 5 (a) shows a few samples of the network power consumption of BeamAdapt, and its upper bound given by the theoretically optimal solution using a brute-force algorithm with client cooperation. The figure also shows the performance of omni directional transmission for comparison. Clearly, the performance of BeamAdapt is very close to the optimal and much better than that of omni. Figure 5 (b) shows the CDF of the additional network power consumption by BeamAdapt compared to its bound: BeamAdapt indeed converges to the optimal solution with a probability of 55%, and only incurs 0.5% additional power compared to the optimal solution when it converges to a sub-optimal.

Using the same network configuration, we can also evaluate the convergence speed of BeamAdapt. Figure 5 (c) shows the PDF of the number of iterations to achieve a small $\varepsilon$, i.e., 0.1% in our simulation. Clearly, BeamAdapt often converges rapidly, i.e., with typically less than three iterations to get a stable SINR.



Samsung Ex. 1024, Page 6 of 12

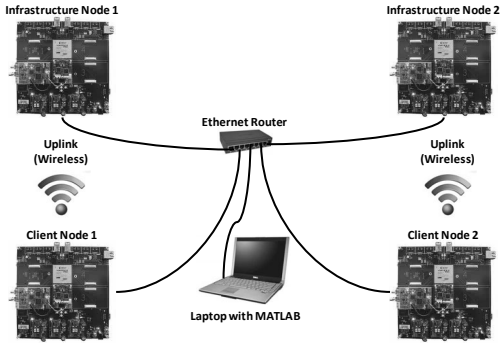


**Figure 6: WARPLab setup for the experimental evaluation of BeamAdapt.**

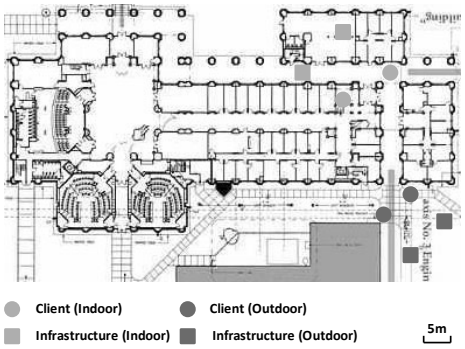


**Figure 7: Environment layout and node locations for the experimental evaluation of BeamAdapt.**

# 5. PROTOTYPE-BASED EVALUATION

In Section 3 we showed that a close-to-maximal beamforming gain can be achieved even when the mobile client rotates at 180°/s. However, compared to static beamforming with a fixed number of active antennas, BeamAdapt faces a new challenge due to its iterative nature: *are mobile clients with BeamAdapt able to timely identify the right number of antennas and transmit power in real-time so that the required SINR is achieved with near-maximal power reduction?* To answer the question, we use WARP to experimentally evaluate the effectiveness of BeamAdapt in realistic environments.

BeamAdapt is compatible with any infrastructure-based network architecture. Since we are not able to conduct experiments on cellular bands due to the lack of license, we instead use ISM band (2.4 GHz) to verify the feasibility of BeamAdapt. Our Qualnet-based evaluation in Section 6 will complementarily show the power saving and network throughput performance of BeamAdapt within a large-scale cellular network.

## 5.1 BeamAdapt Prototype

We realize BeamAdapt using WARPLab, a framework that facilitates rapid prototyping of physical layer designs and algorithms. WARPLab allows symbol-level access to the wireless transceivers embedded on the WARP board, which we leverage to realize the key functionalities of BeamAdapt including beamforming, transmit power and beamforming size adaptation, and SINR measurement. In WARPLab, all WARP nodes are connected through an Ethernet router and a laptop with a MATLAB interface is used to control the nodes, implement the algorithm and collect the measurements. Although our WARP-based prototype does not truly have a power profile for mobile devices, it is our belief that our results will provide the first motivation for wireless modem chipset vendors to seriously consider beamforming for mobile devices.

We have built two types of WARP nodes: one with four antennas implementing BeamAdapt as the client node and the other with a single antenna as the infrastructure node. The physical wireless channel is assumed to be the uplink channel while an Ethernet cable is used to emulate the downlink channel. Since we are only interested in client uplink transmission, we generate dummy frames only at the client node and continuously send them to the infrastructure node.

## 5.2 Experiment Setup

We test the prototype under two physical environments: one inside a building and the other on an empty lawn, both in a university campus. The former and the latter represent typical indoor and outdoor environments, respectively. We use four WARP nodes, including two client nodes and two infrastructure nodes, to form a two-link network. Figure 6 shows our WARPLab setup and Figure 7 shows the locations of the client and infrastructure nodes in the experiment.

While in realistic wireless networks there might be more links that interfere with each other, we consider this two-link network as a reasonable setup for experiments. First, the two-link network is a widely used model in wireless network researches [5], due to its simplicity and generality. Second, even though in realistic such as cellular networks there are more than two base stations within the coverage of a mobile client, the client is often mainly interfering with only one additional base station. This is due to the distributed fashion of placing base stations in a certain area and that each client often connects to the closest base station. Last, since we have selected the ISM bands and the environments of our experiments have continuous but unpredictable wireless transmissions, there are indeed other interference sources at the infrastructure nodes.

In our experiments, we manually add both movement and rotation to the two client nodes. The rotation speed is 0-120°/s, consistent with [1]. The movement is about 0-1 meter per second. Due to the limitation of WARPLab that the WARP boards have to be connected by Ethernet cables, we can only add pedestrian movement speed in the experiment but will simulate a much higher speed in the Qualnet simulation in Section 6.

## 5.3 Findings from Experiments

According to the problem formulation in Section 4, we examine the effectiveness of BeamAdapt in realistic environments with two key metrics, received SINR at the infrastructure node and power consumption by the client node. For received SINR, we examine whether BeamAdapt can closely approach the required SINR even with iteration and client mobility. For power consumption, we compare the power consumption of BeamAdapt with a genie-aided solution which can always correctly pick the right beamforming size and transmit power without iteration. Clearly, the genie-aided solution is always optimal and achieving maximal power reduction. To realize the comparison, we recorded the traces of the channel coefficients during all our measure-



Samsung Ex. 1024, Page 7 of 12

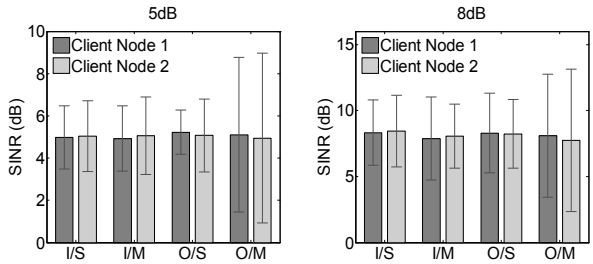


**Figure 8: Received SINR at the infrastructure node in the experiments.**

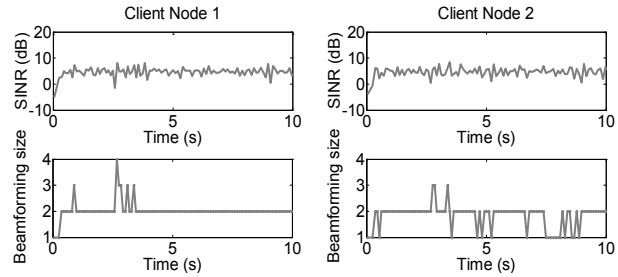


**Figure 9: Received SINR and beamforming size in the experiments at a glance.**

ments and replayed the channel offline to emulate the genie-aided solution.

### *5.3.1 Received SINR*

We first report the received SINR at the infrastructure nodes. To maximally leverage the range of WARP nodes without losing generality, we assume moderate SINR, e.g., 5 dB and 8 dB as the constraint. Figure 8 shows the mean and variance of the received SINR at the two infrastructure nodes, in four scenarios: indoor/static (I/S), indoor/mobile (I/M), outdoor/static (O/S), and outdoor/mobile (O/M).

There are two key observations from the figure. First, BeamAdapt on average can closely approach the required SINR, i.e., 5 dB and 8 dB respectively, for both stationary and mobile client nodes. In most of the scenarios the standard variance is below 3 dB, indicating that the BeamAdapt iteration does not render significant SINR deviation from the target value. Second, in the outdoor/mobile scenario, BeamAdapt yields much higher variance of the received SINR. This is consistent with our observation to the beamforming gain in Section 3.2, due to the lack of compensation by multipath effect to the out-of-date channel estimation and beamforming size.

Figure 9 shows a ten-second snapshot of the received SINR as well as the beamforming size of two client nodes. Clearly, most of the time BeamAdapt is able to timely cope with channel variation and achieve a stable SINR, while the beamforming size is indeed being adapted. We have chosen the measurements in the indoor/mobile scenario to show in the figure while the other scenarios exhibit similar characteristics, as demonstrated by Figure 8.

While BeamAdapt on average achieves the required SINR, it does not guarantee that the SINR is above the target value. This is due to the formulation of BeamAdapt that seeks to use the minimum power to achieve certain capacity. Nonetheless, BeamAdapt will not lead to a large outage probability, since one can simply leave a SINR margin and set the required SINR in BeamAdapt a bit higher than the intended value. For example, if a SINR of 5 dB is needed, one can set 8 dB as the constraint and BeamAdapt will maintain the SINR above the threshold with a probability of 87% according to our measurements.

### *5.3.2 Power Consumption*

We next compare the power consumption of BeamAdapt with that of the genie-aided solution. Again, we note that given the transmit power and beamforming size, the power consumption is calculated using the power model in Section 2.3 instead of measurements. This is because the WARP node uses FPGA boards and programmable RF boards to enable customization, and therefore its power consumption is not optimized and not comparable with realistic beamforming transceivers on mobile devices. Therefore, we use the authentic transmit power but emulate the circuit power to achieve a rational estimation.

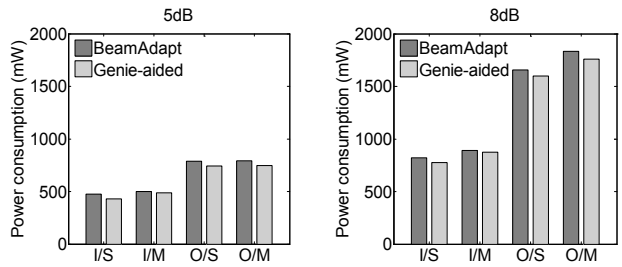


**Figure 10: Power consumption of the client nodes for BeamAdapt and genie-aided solution in the experiments.**

Figure 10 shows the average power consumption of BeamAdapt and the genie-aided solution. Clearly, in all scenarios BeamAdapt closely approaches the theoretically minimum power consumption given by the genie-aided solution, yielding only 5% higher power on average. We note that the genie-aided solution has removed all the imperfections of BeamAdapt in reality, such as converging to a sub-optimal solution, process of iteration, and drop of beamforming gain due to mobility. Therefore, it is the strict upper bound of the power saving performance of BeamAdapt. Not surprisingly, the additional power consumption in our experiments is larger than that in our empirical results in Section 4.5, due to the consideration of all realistic imperfections of BeamAdapt listed above.

## 6. QUALNET-BASED EVALUATION

To complement the prototype-based evaluation, we next use simulation to evaluate BeamAdapt in a large-scale network. To achieve a close-to-reality evaluation, we adopt current cellular protocols and introduce a system design of BeamAdapt that is readily realizable with trivial protocol modification. We employ the simulation tool Qualnet [18] for its open-source feature and support of modern cellular protocols.

### 6.1 Cellular-based System Design

We realize BeamAdapt on mobile clients in a cellular network and again focus on uplink transmission. Due to its distributed nature, BeamAdapt relieves clients in the network from inter-client coordination thereby entails minor protocol modification. There are two key questions one need to answer regarding the system design of BeamAdapt. First, *how does BeamAdapt per-*



Samsung Ex. 1024, Page 8 of 12

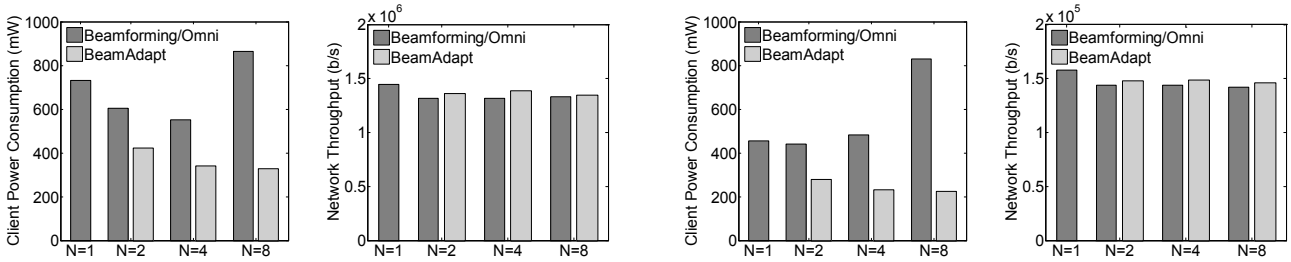


**(a) FTP traffic** **(b) CBR traffic**

**Figure 11: Client power consumption and network throughput comparison between BeamAdapt, static beamforming and omni directional transmission.**

*form uplink CSI estimation?* Second, *how does BeamAdapt obtain the received SINR to perform the beamforming size adaptation?* We next provide the answers.

### *6.1.1 Uplink CSI Estimation*

Due to the absence of uplink/downlink channel reciprocity in cellular networks [19], we can only adopt closed-loop CSI estimation (see Section 2.2) in BeamAdapt. That is, the client concatenates a short field made up of several training symbols to the data field in each uplink frame. Seeing the training symbols, the base station estimates uplink CSI and sends it back to the client. Thanks to the full-duplex property of cellular channels, the estimated CSI can be simultaneously delivered to the client through downlink control signaling while the client is involved in uplink transmission. Therefore, CSI feedback does not incur any additional uplink channel occupation. Moreover, the training field can be very short compared to the entire frame length, i.e., a 16 μs training field for beamforming size of four and a 10 ms frame in UMTS/LTE [19], which further trivialize the overhead of CSI estimation. According to our measurement in Section 3.2, the 10 ms frame length in UMTS/LTE guarantees accurate CSI estimation of BeamAdapt, even with client rotation.

### *6.1.2 Beam Adaptation*

To adapt the beamforming size and transmit power, BeamAdapt needs to know the received SINR of each frame. While it can be similarly sent back to the client through downlink control signaling, we seek to minimize the protocol modification, by leveraging the uplink power control mechanism included in cellular protocols. Uplink power control is widely used in cellular networks to maintain a constant SINR of each client at its base station. It is initiated by the base station, through sending a power control command to the client, containing the value of the required transmit power. Noticeably, this required transmit power is actually $P_O$ in Equation (2), and one can directly identify the optimal transmit power $P_{TX}$ and beamforming size $N$ using $P_O$, as one iteration in the BeamAdapt algorithm. This way, the received SINR is no longer needed by the client and no protocol change is required.

## 6.2 Simulation Setup

Since the beamforming hardware is not included in Qualnet, we have to virtually realize a beamforming system on the client by generating dynamic beamforming patterns in real-time and adopting the power model in Section 2.3 to calculate client power consumption.

We assume the UMTS network system in Qualnet and use the same seven-cell network configuration shown in Section 4.5. However, here we add more clients, i.e., thirty, to mimic realistic base station scheduling and handoff in cellular networks. The area is 4 km×4 km and the base stations have fixed locations, 1.5 km from its neighbors. While the range of each base station is approximately 1 km, we let their coverage overlap similar to realistic cellular networks in urban areas. The clients are allowed to have random linear movement with speed from zero to seventy miles per hour, corresponding to a wide range of client movement speed such as stationary, pedestrian and vehicular. We also incorporate horizontal rotation to the client, with an upper bounded rotation speed of 120°/s, consistent with [1]. We add two applications to the client: FTP with an unlimited-size file to transfer and constant-bit-rate (CBR) with multiple relatively small packets. FTP generates continuous traffic. CBR, on the contrary, creates intermittent traffic by the idle intervals between small-size packets. The FTP traffic has a higher capacity requirement than the CBR traffic.

We evaluate the power reduction benefit of BeamAdapt by comparing it with omni directional transmission and static beamforming with a fixed beamforming size. We examine BeamAdapt and static beamforming with two, four and eight antennas. Note that BeamAdapt with $N$=4 means that the client can select from one to four active antennas (with unused antennas powered off) while static beamforming with $N$=4 always uses four active antennas.

## 6.3 Findings from Simulation

Figure 11 shows the average power consumption of the client as well as the network throughput, under omni, static beamforming and BeamAdapt. We make several key observations. First, BeamAdapt saves more power for the FTP traffic than the CBR traffic since FTP averagely requires higher transmit power. For example, compared to omni directional transmission BeamAdapt with $N$=4 saves 54% and 50% client power for the FTP and CBR traffic, respectively. Second, BeamAdapt with four antennas already provides sufficient power efficiency benefit. The power reduction of BeamAdapt with $N$=8 is only marginally better than that of BeamAdapt with $N$=4. This is due to the confined range of cellular radio signals by the transmit power limitation from clients. Last, the network throughput achieved by BeamAdapt is only slightly lower (<5%) than that by omni directional transmission, and is as good as that by their respective static beamforming counterparts. The slight degradation from omni is due to client mobility and thereby the drop of the beamforming gain, similar to what we observed in Section 3.2.



Samsung Ex. 1024, Page 9 of 12

We also note that the power reduction by BeamAdapt stems from two benefits of beamforming: the reduction of transmit power from the beamforming gain, and the reduction of interference from the directional pattern. Qualnet simulation allows us to further examine the power saving contribution from these two benefits. That is, we first keep the transmit power reduction capability of BeamAdapt only by assuming an omni directional pattern, and then the interference reduction capability only by assuming a beamforming gain of zero. Figure 12 shows their respective contributions to client power reduction, with different distances from the client to the base station. Clearly, as the client moves to cell boundary, i.e., with a larger distance to the base station, both capabilities of BeamAdapt can save more power, and they collectively achieve a higher overall power reduction of the client. This is because when the client is approaching cell boundary, only not the required transmit power increases, but also the interference between adjacent cells is more severe.

# 7. DISCUSSION

We next discuss two important ways to extend BeamAdapt.

## 7.1 BeamAdapt for Receive Beamforming

While we concentrate on transmit beamforming in this work, BeamAdapt can be straightforwardly extended to receive beamforming at the mobile client for downlink performance enhancement. Similarly, we allow a dynamic number of antennas in the beamforming receiver and again the maximal receive beamforming gain is equal to the receive beamforming size. The problem formulation of BeamAdapt in Section 4 still holds only with the client power consumption, $P$, being replaced by

$$P = NP_{Circuit} + P_{Shared}.$$

Because the transmit power is no longer involved, solving the problem is in fact trivial by letting each client use just enough antennas to meet the SINR requirement. Moreover, downlink CSI estimation is even simpler for the client, since the client can directly measure the SINR. As a result, receive BeamAdapt can be easily realized without any modification to the cellular protocol.

## 7.2 Dual Formulation of BeamAdapt

Our problem formulation in Section 4.2 attempts to minimize the network power consumption to achieve certain network capacity. A dual problem to maximize network capacity can be formulated as follows:

$$\text{maximize } C_{Network} = \sum_{i=1}^{M} C_i\,(\boldsymbol{P}_{TX}, \boldsymbol{N})$$

$$\text{s.t. } P_i(P_{TX,i}, N_i) = \rho_i{}', 1 \le N_i \le N_{i,max}.$$

Unlike traditional work that leverages beamforming for maximizing network capacity under a client power constraint, this dual formulation also considers circuit power. As a result, the power tradeoff is still valid and BeamAdapt can be properly modified to provide a distributed solution to this dual problem.

# 8. RELATED WORK

While multi-antenna techniques and directional communication have been generally studied in many other regimes, our work is the first that aims to enable power-efficient real-time beamforming on mobile devices. We next discuss related work in three directions.

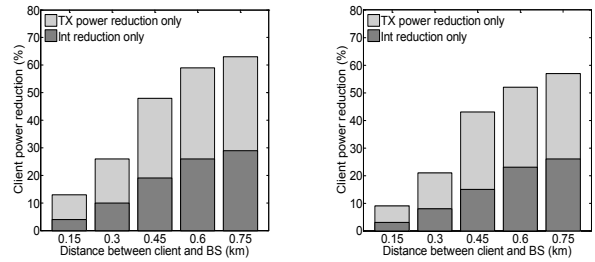


(a) FTP traffic (b) CBR traffic

**Figure 12: Breakdown of client power reduction by BeamAdapt.**

## 8.1 Beamforming

No existing work on beamforming has considered and optimized its use in terms of power efficiency for mobile devices such as tablets and smartphones. Recent work such as [11, 12] considered using beamforming on vehicles to enhance the uplink connection as the client moves. The authors of [20] have experimentally shown the effectiveness of switched beam systems in indoor environments. However, all above solutions use the Phocus Array system [9] and none supports real-time beamforming. More importantly, these solutions do not consider dynamic number of active antennas in beamforming and its power efficiency benefit as we do. Early results from our work on power-efficient beamforming on mobile devices were reported in [21].

## 8.2 Directional Antennas on Mobile Devices

Passive directional antenna is a simple yet inflexible solution to realize directional communication on mobile devices. Many have studied them for infrastructure nodes and mobile nodes that do not rotate, e.g., see [22-29]. Most of the authors focus on MAC protocol designs. In contrast, BeamAdapt is in the physical layer and is complementary to directional MAC designs. Only very recently, the authors of [1, 2] demonstrated the effectiveness of passive directional antennas in improving throughput and power efficiency of mobile devices that can rotate. The solution is based on selecting one out of multiple fixed passive directional antennas. However, there is a key limitation toward their solution: only a limited number of passive antennas are allowed to be implemented, e.g., four in [1], and they are hard to be properly oriented. Such limitation renders a confined gain of their solution due to the failure to cover all directions, i.e., only 3 dB gain using 5 dBi and 8 dBi antennas. In contrast, beamforming with BeamAdapt can easily track channel variation and achieves a guaranteed gain of 6 dB using four antennas.

## 8.3 Energy-Efficient MIMO

While in this work we consider beamforming for its adaptive use in a power-efficient manner, similar concept can be extended to MIMO systems. In [30], we provided a system design of an adaptive MIMO system and experimentally shown that it can minimize the energy per bit of the MIMO transceiver by properly choosing the number of active RF chains. The idea is also explored by the authors of [31] and [32]. The authors of [33] have analytically showed the effectiveness and performance of such adaptive MIMO systems. These solutions, however, are limited to a single link, while BeamAdapt is solving a network problem by optimizing the use of beamforming on multiple mobile clients.



Samsung Ex. 1024, Page 10 of 12

## 9. CONCLUSION

In this work, we reported the first study of beamforming on mobile devices. With both experiments and data from industry, we showed that beamforming is not only feasible but also power-efficient to mobile devices. We then addressed the challenge of identifying the optimal use of beamforming on mobile device, by formulating an optimization problem and providing the BeamAdapt solution. Through both experiments and simulation, we showed that BeamAdapt is able to react to client mobility by promptly identifying the right beamforming size and the transmit power. Collectively it achieves more than 50% power reduction of the clients in a large-scale network.

Client directionality through beamforming is a radical departure from omni directionality assumed by current mobile network paradigms. While we are able to demonstrate its benefit in client power efficiency, more research efforts at various layers of the network system is intended to fully appreciate its potential, which we leave to future work.

## ACKNOWLEDGEMENT

This work was supported in part by NSF awards ECCS/IHCS 0925942, CNS/NeTS-WN 0721894, CNS/CRI 0751173, CNS-0551692, and support from the TI Leadership University program. The authors would like to thank the anonymous reviewers for their useful suggestions.

## REFERENCE

[1] A. Amiri Sani, L. Zhong, and A. Sabharwal, "Directional antenna diversity for mobile devices: characterizations and solutions," in *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2010.

[2] A. Amiri Sani, H. Dumanli, L. Zhong, and A. Sabharwal, "Power-efficient directional wireless communication on small form-factor mobile devices," in *Proc. ACM/IEEE Int. Sym. Low Power Electronics and Design (ISLPED)*, 2010.

[3] L. C. Godara, *Smart Antennas*: CRC Press, 2004.

[4] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly, "Design and experimental evaluation of multi-user beamforming in wireless LANs," in *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2010.

[5] D. Tse and P. Viswanath, *Fundmentals of Wireless Communication*: Cambridge University Press, 2005.

[6] Z. Li, W. Ni, J. Ma, M. Li, D. Ma, D. Zhao, J. Mehta, D. Hartman, X. Wang, K. K. O, and K. Chen, "A Dual-Band CMOS Transceiver for 3G TD-SCDMA," in *IEEE Int. Solid-State Circuits Conference (ISSCC)*, 2007.

[7] T. H. Lee, "The Design of CMOS Radio-Frequency Integrated Circuits," *Cambridge University Press,* 2004.

[8] S. Cui, A. J. Goldsmith, and A. Bahai, "Energy-efficiency of MIMO and cooperative MIMO techniques in sensor networks," *IEEE Journal on Selected Areas in Communications,* 2004.

[9] Fidelity Comtech, *Data Sheet: Phocus Array 3110X* 2008.

[10] D. G. Rahn, M. S. Cavin, F. F. Dai, N. H. W. Fong, R. Griffith, J. Macedo, A. D. Moore, J. W. M. Rogers, and M. Toner, "A fully integrated multiband MIMO WLAN transceiver RFIC," *IEEE Journal of Solid-State Circuits (JSSC),* 2005.

[11] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "MobiSteer: using steerable beam directional antenna for vehicular network access," in *Proc. ACM Int. Conf. Mobile Systems, Applications and Services (MobiSys)*, 2007.

[12] K. Ramachandran, R. Kokku, K. Sundaresan, M. Gruteser, and S. Rangarajan, "R2D2: regulating beam shape and rate as directionality meets diversity," in *Proc. ACM Int. Conf. Mobile Systems, Applications and Services (MobiSys)*, 2009.

[13] WARP, *http://warp.rice.edu/*, 2011.

[14] IEEE International Solid State Circuits Conference.

[15] IEEE Journal of Solid-State Circuits.

[16] G. L. Nemhauser and L. A. Wolsey, *Integer and combinatorial optimization*: Wiley-Interscience, 1988.

[17] R. D. Yates, "A framework for uplink power control in cellular radio systems," *IEEE Journal on Selected Areas in Communications,* 1995.

[18] Scalable Network Technologies, *QualNet Developer: High-fidelity network evaluation software*.

[19] E. Dahlman, S. Parkvall, J. Skold, and P. Beming, *3G Evolution: HSPA and LTE for Mobile Broadband*: Academic Press, 2008.

[20] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, "On the Effectiveness of Switched Beam Antennas in Indoor Environments," in *Proc. Passive and Active Network Measurement (PAM)*, 2008.

[21] H. Yu, L. Zhong, and A. Sabharwal, *Beamsteering on mobile devices: network capacity and client efficiency*: Technical Report 0623-2010, Rice University, 2010.

[22] S. Yi, Y. Pei, and S. Kalyanaraman, "On the capacity improvement of ad hoc wireless networks using directional antennas," in *Proc. ACM Int. Sym. Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2003.

[23] L. Bao and J. J. Garcia-Luna-Aceves, "Transmission scheduling in ad hoc networks with directional antennas," in *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2002.

[24] Y. Ko, V. Shankarkumar, and N. H. Vaidya, "Medium access control protocols using directional antennas in ad hoc networks," in *Proc. IEEE INFOCOM*, 2000.

[25] T. Korakis, G. Jakllari, and L. Tassiulas, "A MAC protocol for full exploitation of directional antennas in ad-hoc wireless networks," in *Proc. ACM Int. Sym. Mobile Ad Hoc Networking & Computing (MobiHoc)*, 2003.

[26] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya, "Using directional antennas for medium access control in ad hoc networks," in *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2002.

[27] M. Takai, J. Martin, A. Ren, and R. Bagrodia, "Directional virtual carrier sensing for directional antennas in mobile ad hoc networks," in *Proc. ACM Int. Sym. Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2002.

[28] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, "DIRC: increasing indoor wireless capacity using directional antennas," in *Proc. ACM SIGCOMM*, 2009.



**Samsung Ex. 1024, Page 11 of 12**

[29] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, "Pushing the envelope of indoor wireless spatial reuse using directional access points and clients," in *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2010.

[30] H. Yu, L. Zhong, and A. Sabharwal, "Adaptive RF chain management for energy-efficient spatial-multiplexing MIMO transmission," in *Proc. ACM/IEEE Int. Sym. Low Power Electronics and Design (ISLPED)*, 2009.

[31] I. Pefkianakis, S.-B. Lee, and S. Lu, *MIPS: MIMO Power Save in 802.11n Wireless Networks*: UCLA Computer Science Department Technical Report TR-100040-2010.

[32] D. Halperin, B. Greensteiny, A. Shethy, and D. Wetherall, "Demystifying 802.11n power consumption," in *Proc. USENIX Int. Conf. Power Aware Computing and Systems (HotPower)*, 2010.

[33] H. Kim, C.-B. Chae, G. D. Veciana, and R. W. Heath, "A cross-layer approach to energy efficiency for adaptive MIMO systems exploiting spare capacity," *IEEE Trans. Wireless Communications,* 2009.

## Appendix: WARP Setup for CSI Estimation

We build a circular array with four antennas on one WARP board as the client node, and use a single antenna on the other WARP board as the infrastructure node. The antenna spacing in the circular array is 0.5 $\lambda$. The client and infrastructure nodes are placed close to the allowed range of WARP board with a moderate SNR (5 dB), i.e., 10 meters in our experiments. The client node continuously sends training symbols to the infrastructure node every 10 ms and the latter sends back the estimated CSI through an Ethernet cable. Therefore, the mobile client updates the CSI every 10 ms, calculates the weight vector and then performs beamforming. To challenge the CSI estimation, we rotate the client node with a computerized motor at 90°/s and 180°/s respectively, while realistic mobile devices rotate at a much slower speed, e.g., 10°/s as the median and 120°/s as the upper bound [1]. We repeat the experiments both indoor and outdoor. While we could not simultaneously examine different beamforming sizes and different CSI estimation frequencies in real time, we have collected traces of the channel coefficients and emulated the channel offline. That is, we replay the channel using the recorded traces but assume different beamforming sizes (2 and 4), and different CSI estimation frequencies (10 ms and 100 ms). Since the beamforming gain is only dependent on the CSI, the offline emulation gives identical results as real-time evaluation does.



Samsung Ex. 1024, Page 12 of 12