

FAST COPPER FOR BROADBAND ACCESS: AN OVERVIEW

Mung Chiang¹, Jianwei Huang¹, Dahai Xu¹, Yung Yi¹, Chee Wei Tan¹, Raphael Cendrillon²

¹ Electrical Engineering Department, Princeton University ² Marvell Semiconductors

ABSTRACT

This is an overview of the ongoing FAST Copper project, which is aimed at substantial improvements in rate, reach, reliability, and quality in copper-last-mile broadband access through fiber/DSL deployment, engineering innovations, and fundamental research. The project is funded by NSF, and is currently pursued jointly by Princeton University, Stanford University, and Fraser Research Lab. In this article, we outline the motivations, challenges, and research issues associated with the project, and report some of the recent results by the Princeton team in each of the four dimensions: Frequency, Amplitude, Space, and Time.

1. INTRODUCTION TO FAST COPPER

FAST Copper is a multi-year, U.S. NSF funded project that started in 2004, and is currently pursued jointly by the research groups of Mung Chiang at Princeton University, John Cioffi at Stanford University, and Alexander Fraser at Fraser Research Lab, and in collaboration with several industrial partners including AT&T. The goal of the FAST Copper Project is to substantially improve the rate, reach, reliability, and quality in copper-last-mile broadband access to everyone with a phone line, through the combination of fiber/DSL deployment, engineering innovations, and fundamental research.

This goal will be achieved through two threads of research: dynamic and joint optimization of resources in Frequency, Amplitude, Space, and Time (thus the name ‘FAST’¹) to overcome the attenuation and crosstalk bottlenecks, and the integration of communication, networking, computation, modeling, and distributed information management for architectural design of broadband access networks.

Access networks are often the rate-reach-reliability-quality bottleneck of end-to-end connections in wide area networks. Realizing the vision of truly broadband and ubiquitous access to almost everyone in the U.S. is a formidable task, with many significant technical and socio-economic challenges. We propose to leverage the installed copper plant, which is by far the most ubiquitous access network in the U.S. The overall solution is a hybrid fiber/DSL deployment where fiber is pushed into the access network but copper takes over the last mile, thereby utilizing the best of ubiquity, broadband, reliability, and economic viability. Can *substantially higher* data rate and application throughput be attained over DSL through research innovations? We believe the answer is definitely positive. To achieve data rates significantly higher than the current levels on low-twist unshielded telephone wires demands thinking about transmission on

copper wires in a new way. This project combines innovative optimization and signal processing techniques with novel network architectures and protocols, as well as an integrated plane of real-time control, computation, data collection, and auto-configuration, to enable an access infrastructure that is *both* broadband and ubiquitous.

After surveying the key ideas in FAST Copper in the next section, we will provide a brief summary of some of the latest developments in 2005-2006 for Princeton’s part of this actively ongoing project. This summary is presented along the dimensions of Frequency, Time, Amplitude, and Space, with more mature results for the Frequency axis of the project. We expect continuous progress to be made by all institutions in the project in the coming years.

2. KEY IDEAS AND TECHNICAL CHALLENGES

Traditionally, DSL broadband access networks have been analyzed by viewing each twisted pair as a separate communication channel, independent of other twisted pairs in the same binder cable, with a fixed pipesize supporting circuit-switched voice traffic. The key to realizing the vision of ubiquitous, readily deployable, and truly broadband access networks is to dynamically optimize the resources in the dimensions of Frequency, Amplitude, Space, and Time, in the multiple-input-multiple-output communication environment of DSL across multiple layers in the protocol stack.

Therefore, there are two shifts of mentality that underlines the wide range of activities in FAST Copper:

- The first key idea is that, instead of holding the traditional view that each twisted pair is an independent channel, we model a bundled cable of twisted pairs as one aggregate *multi-user* communication system. Multiple users *compete* against and *cooperate* with each other in this system. We can explicitly take into account the *crosstalk* effects (both near-end and far-end) that currently form the data rate bottleneck, and to exploit potential *cooperation* in sharing limited resources
- The second key idea is that today’s traffic over broadband access, including voice, data, and video, are predominantly supported by packet switched IP. We can exploit the *burstiness* of the application traffic through aggressive statistical multiplexing, with admission control, traffic shaping, scheduling, and priority queuing mechanisms to ensure the desired tradeoff between the number of application flows supported and the Quality of Service (QoS) attainable.

There are two major bottlenecks to DSL broadband access today: attenuation and crosstalk. We will see solutions from the “Space” dimension of the project to tackle the problem of attenuation, and solutions from the “Frequency”, “Amplitude”, and “Time” dimensions to tackle the problem of crosstalk. Note that we have not even brought in factors such as wider bandwidth and multiple twisted-pairs.

We would like to thank the funding from the NSF CISE ITR program 2004-2008, collaboration with AT&T, and discussions with project co-PI John Cioffi, Alexander Fraser, and their research groups. We also thank the general discussions with Elliot Anshelevich, Marc Moonen, Daniel Palomar, and Wei Yu.

¹FAST Copper project is completely different from the TCP FAST project at Caltech, which is a project that improves the transport layer protocol.

- *Frequency.* In the physical layer, new techniques can be developed based on improving spectral utilization, mitigating multi-user interference, and exploiting multi-user cooperation. Through dynamic adaptation and utilization of frequency spectrum, such as power control, bit loading, or vectored transmission, Dynamic Spectrum Management (DSM) allows maximum flexibility in allocating rates among competing flows, achieves much higher total data rates, and extends the reach of broadband access.
- *Time.* FAST Copper also leverages the potential for time division multiplexing based on the application layer burstiness of data traffic from and to the end hosts. In most communication-theoretic investigations, it is assumed that there is always an infinite backlog of bits that need to be transmitted per user, thus taking out the latency considerations and the temporal dimension. By jointly considering the application layers, burstiness of the required bandwidth provides another degree of flexibility of statistical multiplexing along the temporal axis.
- *Space.* When building robust and efficient broadband access networks, two issues are particularly important: how can a hybrid fiber/twisted pair architecture be designed to utilize the best of fiber-based and copper-based communication potentials, and how can a logical topology be designed to offer fast-recovery after natural failures or malicious attacks?
- *Amplitude.* We propose to install active ‘amplitude control’ mechanisms to shape the flow intensities at the edge to provide different QoS classes through admission control and dynamic bandwidth allocation. At the same time, a network management system constantly probes, measures and monitors the cable and its environments, receives data rate requests from user terminals, and periodically shapes the rate each user is allowed to transmit and receive per time frame.

In summary, by modeling the whole binder of copper wires as one multi-carrier interference channel, with resources ranging from the physical layer to the application layer, we can dynamically optimize over Frequency, Amplitude, Time, and Space, in a stable, robust, and complementary way. Collectively, these four degrees of freedom offer many exciting opportunities to make tangible practical impacts.

3. FREQUENCY: SPECTRUM MANAGEMENT

In this section, we focus on the discussion of Dynamic Spectrum Management (DSM) methods, which are used to mitigate crosstalk (interference) among copper lines in the same binder. The representative DSM algorithms in the literature include *Iterative Water-filling* (IW) algorithm [1], *Optimal Spectrum Balancing* (OSB) algorithm [2], *Iterative Spectrum Balancing* (ISB) algorithm [3, 4], and our recently proposed *Autonomous Spectrum Balancing* (ASB) algorithm [5], all of which are summarized in Table 1. The ASB algorithm overcomes the performance bottleneck of all previous algorithms, and has low complexity, autonomous operation and achieves near-optimal rate region.

We consider a DSL network with a set $\mathcal{N} = \{1, \dots, N\}$ users (i.e., lines, modems) and $\mathcal{K} = \{1, \dots, K\}$ tones (i.e., frequency carriers). Assuming the standard *synchronous* discrete multi-tone (DMT) modulation is applied, transmission can be modeled independently on each tone k as $\mathbf{y}^k = \mathbf{H}^k \mathbf{x}^k + \mathbf{z}^k$. The vector $\mathbf{x}^k \triangleq \{x_n^k, n \in \mathcal{N}\}$ contains transmitted signals on tone k , where x_n^k is the signal transmitted by user n at tone k . Vectors \mathbf{y}^k and \mathbf{z}^k have

Table 1. Comparison of various DSM algorithms

Algorithm	Operation	Complexity	Performance	Reference
IW	Autonomous	$O(KN)$	Sub-optimal	[1]
OSB	Centralized	$O(Ke^N)$	Optimal	[2]
ISB	Centralized	$O(KN^2)$	Near optimal	[3, 4]
ASB	Autonomous	$O(KN)$	Near optimal	[6]

similar structures: \mathbf{y}^k is the vector of received signals on tone k , and \mathbf{z}^k is the vector of additive noise on tone k and contains thermal noise, alien crosstalk and radio frequency interference. We denote the channel gain from transmitter m to receiver n on tone k as $h_{n,m}^k$. We denote the transmit power spectrum density (PSD) $s_n^k \triangleq \mathcal{E} \{|x_n^k|^2\}$, where $\mathcal{E}\{\cdot\}$ denotes expected value. The vector containing the PSD of user n on all tones as $\mathbf{s}_n \triangleq \{s_n^k, k \in \mathcal{K}\}$.

Assume that each user treats interference from other modems as noise. When the number of interfering users is large, the interference can be well approximated by a Gaussian distribution. Under this assumption the achievable bit rate of user n on tone k is $b_n^k \triangleq \log \left(1 + \frac{s_n^k}{\sum_{m \neq n} \alpha_{n,m}^k s_m^k + \sigma_n^k} \right)$, where $\alpha_{n,m}^k = |h_{n,m}^k|^2 / |h_{n,n}^k|^2$ is the normalized crosstalk channel gain (with $\alpha_{n,n}^k \triangleq 0, \forall k, n$), and σ_n^k is the noise power density normalized by the direct channel gain $|h_{n,n}^k|^2$. The bandwidth of each tone is normalized to 1. Each user n is typically subject to a total power constraint P_n , due to the limitations on each modem’s analog frontend: $\sum_{k \in \mathcal{K}} s_n^k \leq P_n$. The data rate on line n is thus $R_n = \sum_{k \in \mathcal{K}} b_n^k$.

The spectrum management problem is defined as follows

$$\max_{\{s_n \geq 0, n \in \mathcal{N}\}} \sum_n w_n R_n \text{ s.t. } \sum_{k \in \mathcal{K}} s_n^k \leq P_n, \forall n. \quad (1)$$

where w_n is a nonnegative weight coefficient of user n . Problem (1) is a tightly coupled (across users and frequencies) and nonconvex (due to interference) problem, thus is difficult to solve even in a centralized fashion.

In order to overcome this difficulty, we introduce the concept of a ‘reference line’, a virtual line that represents a ‘typical’ victim within the DSL system (e.g., the longest line). Then instead of solving (1), each user n tries to maximize a weighted sum of its own rate and the reference line’s rate, i.e.,

$$\max_{s_n \geq 0} w_n R_n + R_{n,ref} \text{ s.t. } \sum_{k \in \mathcal{K}} s_n^k \leq P_n, \quad (2)$$

where the reference line’s rate $R_{n,ref} \triangleq \sum_{k \in \mathcal{K}} \tilde{b}_n^k$, and the achievable bit rate on tone k is defined as $\tilde{b}_n^k \triangleq \log \left(1 + \frac{s_n^k}{\tilde{\alpha}_n^k s_n^k + \tilde{\sigma}_n^k} \right)$. The reference line parameters $\{\tilde{s}^k, \tilde{\sigma}^k, \tilde{\alpha}_n^k, \forall k, n\}$ can be obtained from long-term field measurements. Intuitively, the reference line serves a static pricing term in each user’s optimization problem to avoid purely selfish behavior, and eliminates the need of explicit message passing amongst users. Users then iterate until PSD converge.

The core of the ASB algorithm involves solving Problem (2) for each user n . This is realized by a dual-based decomposition algorithm. First, incorporate the total power constraint into the objective function of (2) using a dual variable λ_n . This decouples problem (2) into one subproblem on each tone k ,

$$\max_{s_n^k \in [0, P_n]} w_n b_n^k + \tilde{b}_n^k - \lambda_n s_n^k, \quad (3)$$

which is still a nonconvex problem but can be easily solved by examining the first order condition. Denote the corresponding solution to be $s_n^k(\lambda_n)$, then the nonnegative dual variable λ_n needs to be updated to satisfy the complementary slackness condition, i.e., $\sum_k s_n^k(\lambda_n) = P_n$ or $\lambda_n = 0$. Although Problem (2) is nonconvex, we know from [2] that the corresponding duality gap of Problem (2) is asymptotically zero (when the number of tones becomes large), thus the dual-based approach leads to an optimal primal solution.

The nonconvexity of the objective function in (3) makes it difficult to prove the convergence of ASB algorithm in general. However, under the high SNR approximation of the reference line (i.e., the reference PSD is much larger than the background noise on all the tones that the reference line is active), we can show the following

Theorem 1 Assume $\max_{m \neq n, k} \alpha_{n,m}^k < \frac{1}{N-1}$, then the ASB algorithm under high SNR approximation globally and geometrically converges to the unique fixed point in an N -user system, with either sequential or parallel updates.

Next we illustrate the performance of the ASB algorithm based on a four-user mixed CO/RT scenario. As depicted in Fig. 1(a), user 1 is the longest line in the network (e.g., a central office (CO) based deployment), whilst the other three users are deployed based on remote terminals (RT). Due to the different distances among the corresponding transmitters and receivers, the RT lines generate strong interferences into the CO line, whilst experiencing very little crosstalk from the CO line. The target rates of users 2 and 3 have both been set to 2 Mbps (by adjusting the corresponding weights w_n). For a variety of different target rates of user 4, user 1 (the CO line) attempts to maximize its own data-rate either by transmitting at full power in IW, or by setting its corresponding weight w_{co} to a very large number in OSB, ISB and ASB. This produces the rate regions shown in Fig. 1(b), which shows that ASB achieves near optimal performance similar as OSB and ISB, and significant gains over IW. For example, with a target rate of 1 Mbps on user 1, the rate on user 4 reaches 7.3 Mbps under ASB algorithm, which is a 121% increase compared with the 3.3 Mbps achieved by IW.




Fig. 1. Performance comparison of various dynamic spectrum management (DSM) algorithms.

4. AMPLITUDE: STATISTICAL MULTIPLEXING

Statistical multiplexing has been widely studied in wireline networks (e.g., [7] and the references therein), where the underlying link capacities are assumed to be fixed. In DSL network, however, the link capacities are closely coupled among users due to crosstalk interferences, and can be dynamically “shuffled” across users based on traffic demands. Here, we aim at joint optimization of various system resources (e.g., rate, buffer) to squeeze as much traffic as possible into the network, subject to flow QoS requirements.

As the first initial step, we focus on the multiplexing of delay insensitive data traffic, where the QoS requirements are defined as the upperbounds of the packet loss probabilities due to buffer overflow. This probability is determined by the traffic characteristics (i.e., mean, peak, variance) and the allocated resource to the traffic (i.e., bandwidth and buffer). When the upperbound is very stringent and the available buffer space is large, the bandwidth requirements of users’ data traffic can be estimated accurately using the concept of *effective bandwidth* [8]. This enable us to model a stochastic traffic with a constant traffic, and perform admission control and statistical multiplexing based on this.

We propose a two-stage Alternate Maximization (AM) algorithm to solve the joint optimization. The first stage optimizes over the rate allocation assuming fixed buffer allocation. The corresponding problem can be simplified and solved using the ASB algorithm in Section 3 in a distributed fashion. The second stage optimizes the buffer allocation assuming fixed rate allocation. The corresponding problem is a quasi-concave maximization problem and can be solved using bi-section search. The algorithm alternates between these two stages until no further improvement can be found. Since each stage of the AM algorithm improves the system objective, which is upper-bounded, the AM algorithm always converges. Details of the algorithm can be found in [9].

5. SPACE: ACCESS TOPOLOGY DESIGN

The “space” dimension of FAST Copper consists of two types of problems: architectural decision problems, and topology design problems. The former considers with the division of functionalities between access and core networks, e.g., how large the access networks should be, where the various types of video servers should be placed, which network elements should be responsible for mitigating crosstalk or for reducing excessive bandwidth demand, etc. Here we will focus on the more tactical problems of survivable topology design.

A main target of topology design in FAST project is to offer fast recovery for access networks after natural failures or malicious attacks. The access parts of the network infrastructure aggregates increasing volumes of voice, data, and video traffic from end users, but are usually the least protected, in contrast to ring-based metro networks and partial mesh backbone networks. This lack of protection makes the access network the bottleneck of end-to-end survivability, and is because of the economic reason of high price per customer when provisioning survivable networks. It is necessary to design survivable tree topologies, through the appropriate addition of a limited number of redundant links to the tree, that can provide the best survivability-cost tradeoff.

The structure of the access network is a “fat” tree rooted at central office, i.e., for an intermediate terminal node within an access network, the capacity/traffic of its upstream link is the aggregation of the capacity/traffic of all its downstream links. Therefore, to recover from its upstream link failure, the terminal has to relay the traffic from another terminal of the same or higher level. Such feature of the access network makes the problem of designing reliable access network different from that of designing reliable backbone network.

Different optimization problems can be derived from the variations in the existence of fat tree, single level or multi-level tree, objective function (minimizing total cost to construct survivable access network or maximizing total revenue with limited budget) and the cost function of edge capacity (concave or uncapacitated fixed cost). We outline two new problems as follows.

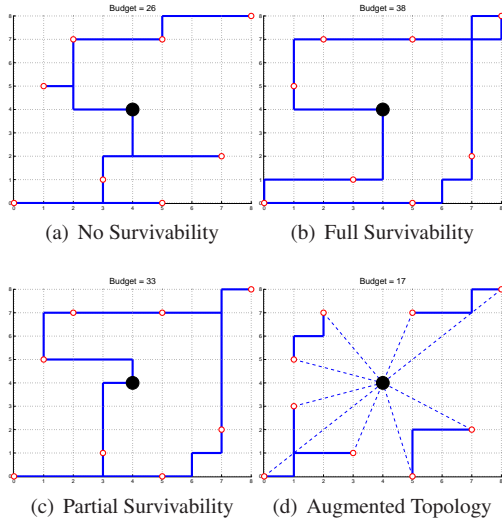


Fig. 2. Access network design for grid network

Budget-Constrained Revenue Maximization

In this problem, uncapacitated fixed cost model is adopted and there is no existing tree. With a limited budget in network construction, we need to search for a subgraph to maximize the total revenue from providing survivability for partial remote terminals. The problem can be proved to be NP-hard through the reduction of Steiner Tree problem.

In the sample Manhattan-like grid network shown in Figures 2(a)~2(c), the empty circles are remote terminals and the solid circle at the center is central office. Each edge has 1 unit cost and the revenue by providing survivability (2-connectivity) to each remote terminal is 1 unit (different from unit edge cost). Fig. 2(a) shows the minimum cost (26 units) access network design without any survivability, and Fig. 2(b) shows the minimum cost (38 units) access network design with full survivability as well as the corresponding 8-unit revenue. In contrast, Fig. 2(c) shows the maximum revenue is 6 units when the budget is constrained by 33 units.

Provisioning Survivability for Existing Single-Level Tree

In this problem, we search for the minimum cost incremental topology design to provide full survivability to all the remote terminals within an existing single-level access tree, which connects the central office and each remote terminal directly. Uncapacitated fixed cost model is used. This problem is poly-time solvable, and a sample minimum cost augmented graph is shown as bold lines in Fig. 2(d) where the existing access tree is shown with dotted lines.

6. TIME: MULTIUSER SCHEDULING

So far we have been focusing on throughput in terms of rate region and multiplexing gain. However, there are also other QoS metrics important to triple play traffic, including delay, jitter, and fairness. These considerations in part determine which point on the rate region should the system attain, and require a finer granularity of control among the admitted flows. We also need to incorporate the stochastic dynamics of flow arrival and departure.

At each user, dynamic bandwidth or buffer allocation among flows that have different QoS requirements is necessary. Table 2 il-

Table 2. Different QoS requirements and control mechanisms.

QoS requirement	Characteristic	Control mechanism
Average throughput	Statistical	Multiuser scheduler
Average delay	Statistical	Multiuser scheduler
Inter-user fairness	Deterministic	Multiuser scheduler & Adm. Ctrl.
Packet loss	Statistical	Priority queueing & Adm. Ctrl.
Hard delay bound	Deterministic	Priority queueing
Jitter	Statistical	Priority queueing

lustrates how different QoS requirements from each traffic flow may be met by the appropriate control mechanisms. At the inter-user level, average throughput, average delay and fairness are ensured by the multiuser scheduler. At the intra-user level, priority queueing is used to differentiate between real-time and non-real-time flows.

7. CONCLUSION

FAST Copper projects presents new networking research challenges, motivates a wide range of difficult problems in fundamental research disciplines, and offers an opportunity for research to make tangible impacts on how people access information.

There are two types of research challenges in FAST Copper project. One is the array of problem formulations in the individual dimensions of Frequency, Amplitude, Space, and Time. These problems range from communication theoretic ones to graph theoretic ones, and from nonconvex optimization to stochastic systems. The other is the need to quantify architectural principles in broadband access networks. This requires a characterization of horizontal and vertical decompositions for functionality allocation. Both types of questions also have applicabilities beyond the scope of FAST Copper project itself.

8. REFERENCES

- [1] W. Yu, G. Ginis, and J. Cioffi, "Distributed multiuser power control for digital subscriber lines," *IEEE Journal on Selected Areas in Communication*, vol. 20, no. 5, pp. 1105–1115, June 2002.
- [2] R. Cendrillon, W. Yu, M. Moonen, J. Verlinden, and T. Bostoen, "Optimal multiuser spectrum balancing for digital subscriber lines," *IEEE Transactions on Communications*, vol. 54, no. 5, pp. 922–933, May 2006.
- [3] R. Cendrillon and M. Moonen, "Iterative spectrum balancing for digital subscriber lines," in *IEEE ICC*, 2005.
- [4] R. Lui and W. Yu, "Low-complexity near-optimal spectrum balancing for digital subscriber lines," *IEEE ICC*, 2005.
- [5] J. Huang, R. Cendrillon, M. Chiang, and M. Moonen, "Autonomous spectrum balancing (ASB) for frequency selective interference channels," in *IEEE ISIT*, 2006.
- [6] R. Cendrillon, J. Huang, M. Chiang, and M. Moonen, "Autonomous spectrum balancing for digital subscriber lines," submitted to *IEEE Transactions on Signal Processing*, 2006.
- [7] M. Reisslein, K. Ross, and S. Rajagopal, "A framework for guaranteeing statistical QoS," *IEEE/ACM Transactions on Networking*, vol. 10, no. 1, pp. 27–42, 2002.
- [8] F. Kelly, "Notes on effective bandwidths," *Stochastic Networks: Theory and Applications*, vol. 4, pp. 141–168, 1996.
- [9] J. Huang, C.-W. Tan, M. Chiang, and R. Cendrillon, "Statistical Multiplexing over DSL Networks," submitted to *INFOCOM* 2007.