ArrayComm

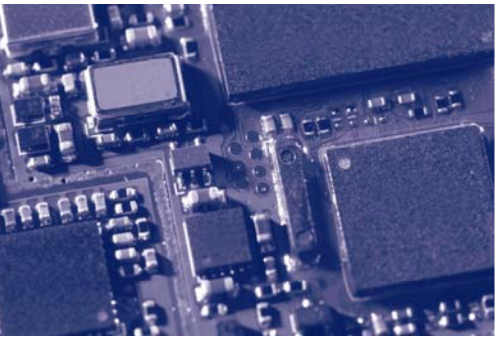

Multi-Antenna Signal Processing: Drivers of Adoption in 3.5G, 3GLTE, and WiMAX

October, 2006

A YGOMI COMPANY

Introduction

Wireless operators are increasing their focus on data and multimedia services to drive revenue growth. As we'll describe in more detail here, this is creating demands for substantially improved radio equipment performance. Unfortunately, years of innovation in wireless have left little new technology ore to be mined for performance improvements — with the exception of the dimension of *space*. Multi-antenna signal processing (MAS) software provides more control over the spatial distribution of radio energy, yielding well-proven order-of-magnitude performance improvements. As a result, MAS is now being embraced as a key part of next generation wireless networks.

MAS technology (also known as *smart antennas*, *space-time processing*, or *MIMO* — see our note at right on terminology confusion in the industry) has been discussed on in many venues throughout the industry, with a primary focus on the technology itself. Here we hope to provide more commercial context for the drivers of adoption in 3.5G, 3G-LTE, and WiMAX standards, summing up with a set of buyers' considerations for those of you employing the technology in your client devices, infrastructure equipment, or networks.

What's MAS?

There is a lot of confusion in the industry today about multi-antenna architectures and processing modes. Commonly-used terms such as *smart antennas*, *MIMO*, *adaptive antennas*, or *beamforming* mean many different things to different people. The more general term *multi-antenna signal processing* (MAS) best fits the software category we're discussing in this article — providing a more all-encompassing and generic label for the complete taxonomy of multi-antenna approaches.

The Keys to Revenue Growth for Operators

Wireless network operators are pursuing new sources of revenue growth as their current voice service markets become saturated, and as competition pushes down voice ARPU despite rising MOU. Non-voice service menus now go beyond just ringtones and SMS to include mobile video, Internet access, and myriad new applications (from on-line gambling to location-based traffic updates and m-commerce wallet functions).

The revenue share of data services in Asian markets is already on firm footing (DoCoMo in Japan[1] and SKT in Korea[2] both receive 27% of ARPU from non-voice services today), and the rest of the world is catching up (19% of ARPU for Vodafone in Europe, and 11–12% for the major US operators[3] [4]). All indications point to continued solid growth ahead.

Different Fundamentals

Conversation about services beyond voice often glosses over an important point. The fundamentals of subscriber economics and experience metrics for data and video are *substantially* different from those for voice.

Figure 1 tells a stark story about the economics of data and video. Using current mass-market prices in the US to indicate subscriber value for voice, data, and video services, dividing by the capacity they consume on average for each of these media every month yields a dramatic illustration of the differences in willingness to pay per unit of capacity consumed. The conclusion: voice, data, and video services are worth roughly $1.00, $0.10, and 0.3¢ per MB to subscribers, respectively! (For reference, on-demand movies alone net out to about 0.9¢ per MB.) A wireless cost structure that supports voice will require *immense* changes in the long run to support mass-market data and multimedia services profitably.

And what about a "mobility premium" for these wireless services? Note that cellular voice didn't reach mainstream adoption till its prices approached those of wired telephony. Further, leading indicator networks in Australia[5] show very minimal premiums for truly mobile broadband access in practice. It appears that in the long run the service premium for mobility is small.

Subscriber Willingness to Pay by Service per Unit of Capacity Consumed

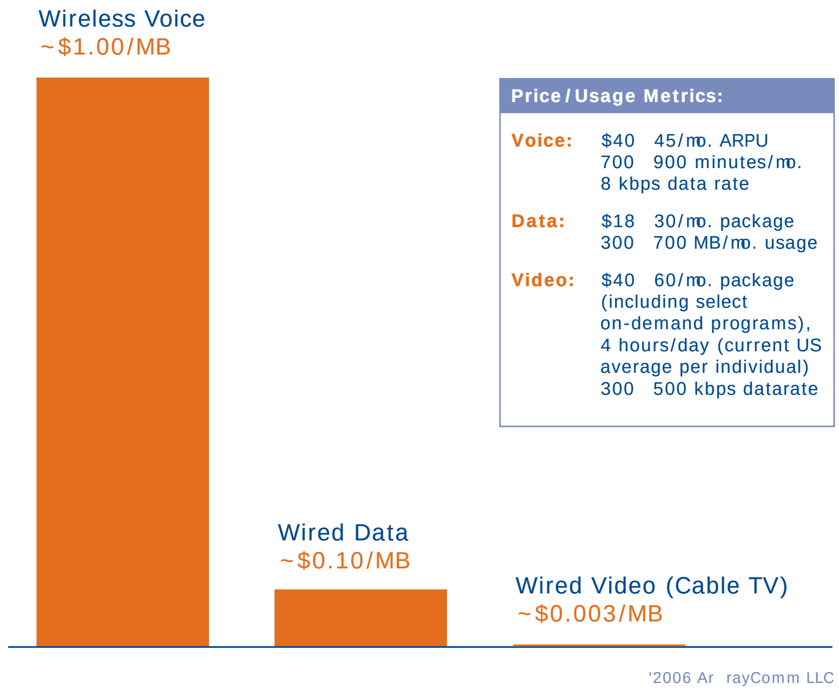

Figure 1

When expressed as willingness to pay per unit of network capacity consumed, subscriber behavior for voice, video, and data services show dramatic differences in fundamental economics[6]

The other fundamental difference for non-voice services is client data rate. With voice, it is very difficult for a subscriber to see or hear performance beyond the largely binary feedback of "has my call been dropped or not?" In contrast, with high-bandwidth applications like broadband Internet access or mobile video, the new dimension of client data rate becomes immediately obvious to users. They can watch their download or upload data rate — or video frame rate and quality — climb as they approach a base station in the network and then slow to a crawl as they reach a point of minimum signal and maximum interference at the cell edge. Product or service reviewers in the press can do their own thorough performance tests, and credible word-of-mouth reports on this performance metric already spread quickly on the Internet. This spells new stress for operators and manufacturers concerned about share positions and brand assets.

Avenues for Performance Improvement

The question now is, how can >100x performance improvements be achieved? We note here a number of avenues commonly considered:

Area	Scale of New Gains Likely
Shannon's Law. Performance of current wireless systems is very close to that predicted by Shannon's Law that defines error probability in noisy communication channels.	**~0**
Time. The efficiency of MAC-layer assignment of users to resources in current wireless gear is also already well advanced.	**~0**
Time and Frequency. W-CDMA and OFDM (which, despite much industry buzz about the latter, are roughly equivalent in the time-frequency domain) already provide very sophisticated organization of users in time and frequency, so there is little additional progress to be made here.	**~0**
Frequency Band Sharing. Cognitive radio — which allows subscribers to 'squat' in unused spectrum on a not-to-interfere basis — is only in early research stages. Even when it becomes practical, there will be little real gain in any market of substantial commercial interest. Mobile broadband spectrum will be very heavily used, leaving little room for cognitive-radio squatters to sneak in.	**maybe ~2x, someday**
Frequency Expansion. Instead of increasing efficiency, one can always just use more spectrum. Unfortunately the spectrum most appropriate for mobile applications — between 500 MHz and 3 GHz — is already well utilized by military, public safety, and commercial operations in markets of any interest. Although there are smaller blocks available, it is no longer possible to achieve anything like an order-of-magnitude gain in the industry-aggregate availability of spectrum for mobile applications.	**~2x**

Area	Scale of New Gains Likely
Economies of Scale. One can also simply reduce the unit cost of the radio equipment itself. The strategy behind broad initiatives like WiMAX includes bringing high economies of manufacturing scale to bear on cellular wireless. Given requirements for carrier-class reliability, radio equipment for widely-adopted cellular standards already enjoying reasonable manufacturing scale, and site costs other than the radio equipment dominating cost structures anyway — scale will make a meaningful contribution, but it can't solve the problem all by itself	**~2 to 4x**
Space. There is plenty of ore left to be mined in the vein of radio system design by more fully utilizing the dimension of space — in fact, at least 10x in the immediate future and a lot more in the long term. And the mining of space need not wait for new innovation. It requires merely vigorous application of MAS technology that's already well proven.	**>10x**

Tapping the Space Dimension

Current wireless networks employ comparatively blunt instruments for the distribution of radio energy in physical space (see Figure 2). As this greatly simplified signal-pattern diagram suggests, this approach creates vast amounts of waste in the system. Power is distributed where subscribers aren't and self-interference is created that degrades signal quality.

Figure 2

Multi-antenna signal processing (MAS) enables much tighter control over the distribution of radio energy in space.

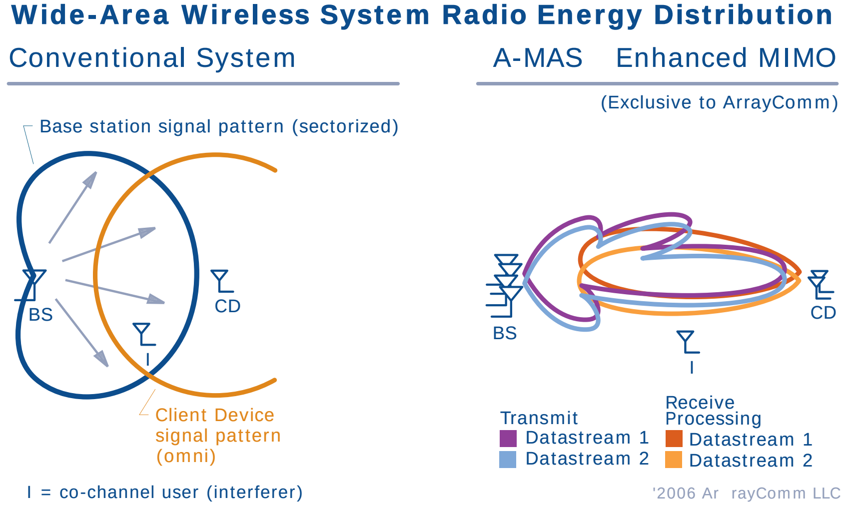

An approach using MAS software, in contrast, takes precise control of the space dimension and puts radio energy only where it's really required (see Figure 2 at right). MAS software drives an array of two or more antennas on either the client device, the base station, or both, leveraging the principle of coherent combinations of radio waves to create a focus of transmit energy (or receive sensitivity) on the intended recipient (sender) and the absence of energy (sensitivity) on sources of co-channel interference. MAS-enabled devices can take advantage of a number of possible gains from using multiple antennas, including link budget improvements from both diversity and combining gains, along with client data rate and overall network capacity benefits from active interference mitigation and spatial multiplexing.

Figure 3 below illustrates how these different categories of MAS benefit add up to improve the overall performance envelope for wireless systems.

Figure 3

MAS software drives significant improvements in wireless device and system range, client data rates, and capacity.

MAS Impact on Wireless System Performance

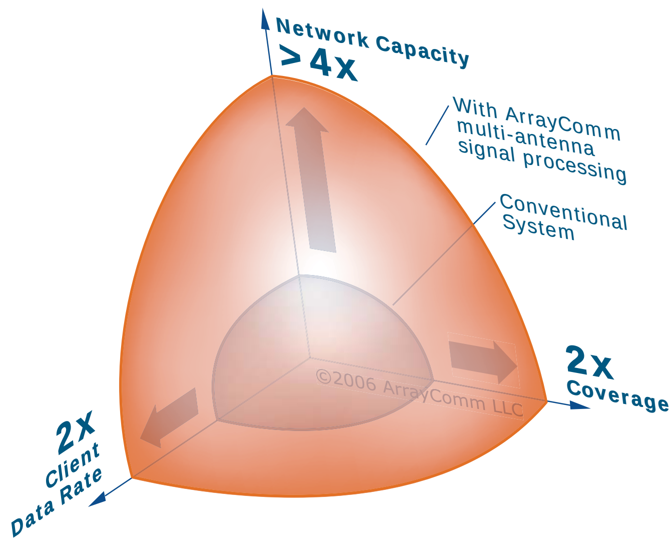

The Industry is Embracing MAS for Next-Generation Systems

In response to operator's demands for performance improvement, MAS takes different forms in existing networks and in next generation mobile broadband networks.

Client Device Solutions in 3.5G Networks

Operators rolling out services based on HSDPA and EV-DO technology are discovering the client data rate problem outlined above. They are finding that their existing 3G network footprint creates "Swiss cheese" coverage for high-data-rate services, and that their suppliers' client devices' performance varies widely from one model to the next.

Because there are often significant constraints on deploying MAS-based infrastructure solutions in existing networks, 3.5G operators are urging their client device suppliers to develop MAS implementations on the client side. In response, Qualcomm has added receive diversity and some simple coherent-gain processing to their chipsets, and EV-DO data cards from Sierra (for example) have implemented two-antenna solutions on this basis. Others are working to catch up on diversity solutions as well. To indicate what's coming next in this area, we show in Figure 4 the output of detailed real-world network simulation (using a European operator's planning tool and network configuration) for ArrayComm solutions that add interference mitigation to the diversity+combining baseline established by Qualcomm. Even in lightly-loaded networks, interference from the neighboring base station or a single dominant co-channel user can substantially degrade data rates at cell edges. Adding a second antenna to the client allows it to ignore only one co-channel interferer, but the net result is significant — more than doubling data rates throughout the network in real-world conditions.

Figure 4

MAS-enabled client devices in HSPA networks deliver more than 2x improvement in subscriber experience.

MAS Impact on HSPA Subscriber Experience

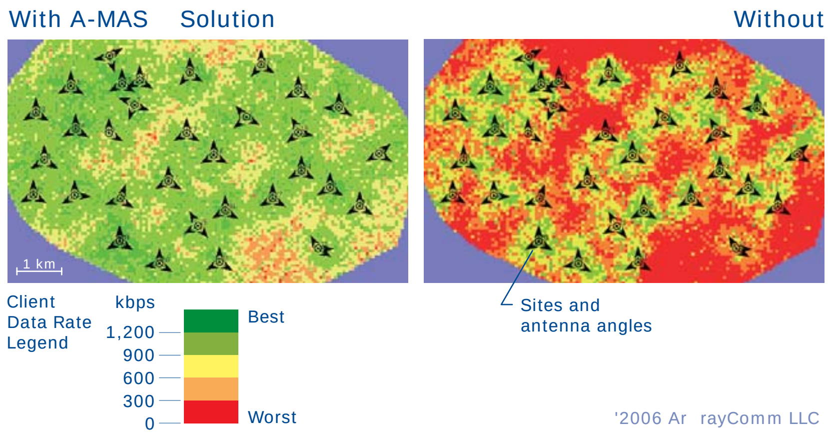

MAS in Next-Generation Mobile Broadband Networks

The 3GPP discussions of 3G LTE (Long Term Evolution) and the Mobile WiMAX profiles established in the WiMAX Forum process are both embracing the combination of OFDM and MAS architectures. Driven in large part by operators who see successful realization of MAS' promise as essential to their own strategic success, these discussions are motivating a renewed supply-side interest in implementing MAS architectures on both the infrastructure and client sides of the equation.

A Brief Buyer's Guide: Not All MAS Approaches Are Created Equal

Many approaches to MAS concepts have been attempted over the past 15 years. Some early trials involved large, expensive, and precisely-calibrated arrays that in the end didn't work very well, and some involved so-called “appliqué” solutions — aftermarket add-on boxes that also yielded generally poor performance because of limited integration with the existing radio hardware and necessarily unsophisticated algorithms.

As the MAS software source with 99% unit share of commercial wide-area deployments of the technology[7], and given our unmatched 14 years of experience in the field, ArrayComm is well-equipped to offer a few guidelines for MAS implementation and to help the industry build on our foundation of lessons learned. A technical discussion of specific algorithms is beyond the scope of this article, but here are some general principles:

Do your homework thoroughly. Many tools for network or economics analysis and performance simulation from single-antenna domains (e.g. interference averaging) yield misleading results when applied to MAS-enabled gear. Getting MAS analysis right is admittedly more complicated, but essential.

Think integrated. The highest performance for the least marginal unit cost is obtained by integrating MAS into client and infrastructure designs from the outset, not adding them on after the fact.

Consider network performance, not just the link. MAS modes that achieve useful results for an individual link (e.g. the baseline form of STC MIMO in WiMAX) can fail in a multi-cell, multi-user context. Network-level analysis of fully-loaded systems is essential.

Use multiple approaches. Use the right tool for the job. Operator requirements and subscriber behaviors vary, from one market to another and from one moment to another. Different MAS architectures have unique strengths and weakness in different applications — there is no single "best" approach. It is much better to include all approaches in the system (possible, since this is software) and let environmental conditions dictate which one is used. This brings us to...

Anticipate dynamic, seamless use of all approaches. We have shown in our PHS implementation, where eight different MAS algorithms are selected on the fly for optimized performance — on a frame-by-frame and user-by-user basis — that MAS architectures can be very dynamic systems. There are many levels of radio system control (beyond individual cells to the network level, for example) that can be incorporated into this sort of self-organizing optimization process. Stay tuned for further innovation in this area!

In Summary

Wireless operators face a sizeable technology challenge as they pursue growth through data and multimedia services. We have outlined how MAS software will contribute to meeting that challenge, along with continued spectrum efforts and new scale economies in the wireless supply base. The wireless future the operators envision for their subscribers is an exciting one — you can expect to see MAS play a key role in making it happen.

End Notes

[1] NTT DoCoMo, Inc. Annual Operating Data (see http://www.nttdocomo.co.jp/english/corporate/investor_relations/business/fiscal_e.html)

[2] Unstrung, "CDMA Reports on B'band", May 8, 2006

[3] Cingular Wireless press release, "Cingular Wireless Reports Strong Second-Quarter 2006 Results", July 20, 2006

[4] Unstrung, ibid.

[5] Personal Broadband Australia and Unwired.

[6] Based on latest US-market wireless voice ARPU and MOU data from Cingular and T-Mobile (company materials); consumer packages for wired broadband Internet access from AT&T (ADSL) and Comcast (cable modem); typical broadband Internet usage figures from Earthlink; consumer cable TV packages from Comcast; and the latest A.C. Nielsen figures for average amount of television watched per day per individual in the US (a mind-boggling four hours).

[7] Visant Strategies, "Intelligent Antennas", May 2006.