A Conversation with Teresa Meng

PHOTOGRAPHY BY TOM UPTON

In 1999, Teresa Meng took a leave of absence from Stanford University and with colleagues from Stanford and the University of California, Berkeley, founded Atheros Communications to develop and deliver the core technology for wireless communication systems. Using a combination of signal processing and CMOS RF technology, Atheros came up with a pioneering 5 GHz wireless LAN chipset found in most 802.11a/b/g products, and continues to extend its market as wireless communications evolve.

As a result of this effort, Meng was named one of the

The founder of Atheros ANALYZES THE
ROLE OF SIGNAL PROCESS-
ING IN THE EVOLVING
WORLD OF WIRELESS
COMMUNICATIONS.

Top 10 Entrepreneurs by Red Herring, Innovator of the Year by MIT Sloan School eBA, and received the CIO 20/20 Vision Award.

Now, as Atheros prepares for its proposed IPO, offering up to \$100 million of common stock, Meng remains on the board of directors, but has returned to Stanford and has been appointed the Reid Dennis Professor of electrical engineering to con-

tinue her research there. Her current focus is on circuit optimization, neural signal processing, and computation architectures for future scaled CMOS technology.

Born and raised in Taiwan, Meng graduated from National Taiwan University before coming to UC Berkeley, where she earned her Ph.D. in electrical engineering and computer sciences. She joined Stanford in 1988.

Conducting the interview with Meng is Robert Broder-son of UC Berkeley, professor in the department of electrical engineering and computer sciences. He served as an advisor to Atheros when Meng founded the company.

Broder-son received his Ph.D. from MIT in 1972 and then spent three years with the Central Research Laboratory at Texas Instruments before joining the faculty at UC Berkeley. His research has involved signal processing applications and new applications of integrated circuits,

focused in the areas of low-power design and wireless communications.

In 1998 he founded the Berkeley Wireless Research Center, a research effort involving the design of highly integrated MOS wireless systems. Broder-son is an IEEE fellow and the winner of numerous awards for his research.

ROBERT BRODERSON: Could you talk about what Atheros does?

TERESA MENG: Atheros as a company has endeavored to be the technology leader in wireless systems design—architecture, circuit design, software, system performance, and all that. When we talk about wireless today most of us think about cellphones—which use licensed band and deliver low data rates and very limited service capability. Atheros was founded to provide the technology to change

that view so that when people think about wireless in the future, they won't feel its presence, as there will be almost unbounded, unlimited capacity everywhere, and it will be much cheaper and easier to use.

■
RB: I would assume the substantial part of this is a signal processing technology, but could you be more specific on what you mean by that technology?

TM: There are three basic ingredients in the technology. The first is definitely signal processing. In the past two decades a lot of research and industry achievements have made it possible for us to understand how to transmit signals through a wireless medium, based on sophisticated signal processing algorithms.

Second, in the late '90s, it became possible to imple-

ment gigahertz RF circuits using digital CMOS technology—the predominant technology that people have been using for implementing microprocessors and memory.

The third ingredient is the opening up of the unlicensed band. Before 1997, for example, a carrier would have to pay a lot of money—in the billion-dollar range—to the government to have the right to use a bandwidth of several megahertz for their cellphone service. In 1997, the U.S. government opened up the 5GHz UNII band that allows unlicensed users—everybody in the United States—to use up to 550 megahertz of bandwidth, as long as they follow the rules.

With the availability of wide bandwidth and CMOS technology being advanced enough to process the bandwidth at this frequency, and with the signal processing know-how—all this created what I call the “wireless

revolution” that freed us from the previous notion that wireless communication is expensive, inherently constrained by a low data-rate, and is scarce. Bandwidth used to be a very scarce commodity, which is not true anymore with the opening up of unlicensed bands. This is the path Atheros would like to lead: to change people’s view of wireless service from tele-communication to more of a data-communication notion where equipment can be updated very quickly and inexpensively, and basically provide a level field for competition.

RB: People usually call this Wi-Fi [wireless fidelity].

TM: Wi-Fi was originally set up to provide the interoperability among the 802.11b devices that we had before. Wi-Fi has expanded to embrace almost all the wireless LAN devices being used in the market today, which include not only 2.4 GHz devices but also devices operated in 5 GHz. I expect this notion will expand to include ultra-wideband devices and, in the future, 60 GHz radios. Even though these devices will be built by different companies, they will interoperate so that users can have a choice of the best and the cheapest product. The Wi-Fi notion provides the rules for a democracy in the wireless communication industry in that it lets everyone participate.

RB: What were the techniques you used to do the signal processing design at Atheros? What was unique about what you did?

TM: The reason we started Atheros was that we had a chance to do things right so that the benefit of the technology could be enjoyed by the great majority of the people. What we decided to do was to throw away all the conventional wisdom of how to design a radio and look at the application requirements for delivering the best possible performance, which is the theory part, and try to achieve that performance using very practical design techniques.

We then researched and developed the architectures and algorithms that are most suitable for efficient CMOS implementation. As a result, Atheros was able to then deliver wireless transceiver chipsets, now in its fourth generation, very quickly and successfully, almost always working on the first silicon. To deliver this kind of robustness in design, we had to rely on extensive simulation. Every piece of silicon we ever built has a simulation model in the system performance evaluation tool that we developed in the company during the first year.

For example, analog components have many impairments—phase noise or nonlinear distortion—which need to be modeled accurately. This, in combination with the

digital signal processing done on the received signal after A/D [analog to digital conversion] and on the transmitted signal before the D/A [digital to analog conversion], modeled down to the exact bit-true level, allows us to accurately predict the performance of the final product long before any circuit design starts.

I think this practice probably laid a very good foundation for generations of products to be implemented in a much more reliable way than what can be supported by other design methodologies.

RB: There are many different ways to implement signal processing. For example, Texas Instruments has signal processing chips. Can you describe how your approach is the same or different from what others do?

TM: Definitely, it’s quite different. *DSP* actually has two definitions. First, it’s *digital signal processing*, and that’s what I refer to. The second definition is *digital signal processors*. That’s a certain kind of programmable processors that were designed specifically for running DSP-like applications.

The problem with using DSP processors in our business is that the amount of computation required in today’s wireless communication design is simply too high for a programmable processor. For example, the amount of signal processing required in an 802.11a/g baseband transceiver is 50 GOPS [giga, or billion, operations per second]. Today’s DSP processors can probably deliver several hundred MIPS [million instructions per second], or we just say one GOPS, per processor. That means, in order to deliver the kind of performance that we require in our chipset, we would need at least 50 of TI’s processors—and these are very high-performance processors. Not to mention the cost, and its unrealistic power consumption. Using dedicated hardware designed for signal processing algorithms, we can achieve a power efficiency of approximately 1 GOPS per milliwatt. This means that to deliver an 802.11a/g-type system, we need to consume only about 50 milliwatts for signal processing. Imagine using 50 of TI’s processors; it’s definitely not something you can carry around within your computer. Therefore the power consumption requirements alone have made the programmable solution completely impractical.

RB: It seems like one of the areas where we are going in the future is lots of radios that are able to communicate through many different kinds of standards. Doesn’t that require some sort of programmability?

TM: Flexibility is different from programmability. Flexibility can be obtained through hardware as well. For

example, even though we designed dedicated hardware silicon, there are at least several hundred parameters that can be modified during the operation of the wireless transmitter/receiver pair (or radio). So multistandard, multimode radios will use dedicated hardware, but with flexibility built in. It's a different design approach. There is no way that any programmable processor in the near future can deliver the kind of power performance at the cost that's required in the future wireless applications.

RB: Where is the wireless world going? What do you think are going to be the big future applications of DSP for wireless?

TM: It's interesting that in the wireless field the actual accumulation of knowledge and development far exceeds what we can implement today. With the opening up of the unlicensed bands, the most striking fact, which most people have not realized, is that bandwidth is unlimited. Most people started with this 100-year-old assumption that wireless bandwidth is very scarce and therefore you have to use it very sparingly. That assumption is completely invalid today.

RB: Why is it invalid? There is only so much spectrum available, right?

TM: In the 5 GHz band, we have 550 megahertz of bandwidth. In ultra-wideband, we have seven gigahertz of bandwidth, from 3 to 10 GHz. And in the 60 GHz carrier band, we have another seven gigahertz of bandwidth. So the bandwidth is unlimited in the sense that we haven't figured out how to use it appropriately. The technology has not caught up with the fact that we actually have 14 or 15 gigahertz of bandwidth at our disposal. Until we have a grip on how to use this bandwidth, the amount of capacity available to us is essentially unlimited.

So far, a 3G cellphone base station has a capacity of about 2 megabits-per-second maximum to be shared by 50 or 60 people. But with 14 or 15 gigahertz of bandwidth, we can build a base station of 100 gigabits per second for people to share, and on top of that, we can exercise spatial diversity, frequency diversity, and time diversity. The amount of diversity and the degree of freedom is so large that the computation or signal processing required to leverage this large amount of bandwidth is phenomenal.

We need to capitalize on the unlicensed band capability. The amount of signal processing required to do that is another order of magnitude higher than what we are capable of today. As silicon continues to scale down, the power efficiency and computation capability will

increase, and that's exactly what we need to implement these very interesting signal processing algorithms so that another level of service can be provided in an extremely easy and low-cost way.

RB: Do you see ultra-wideband as being one of the future technologies?

TM: I think ultra-wideband is very interesting in that it has seven gigahertz of bandwidth. The problem with ultra-wideband as it stands now is that the total transmit power allowance is less than 1 milliwatt. This severely limits the range of these radios. In the 5 GHz or 60 GHz bands, the power limit is much higher—at least 100 times higher. Therefore I view ultra-wideband today more like a playroom for the industry and academia to practice and learn how to design truly wideband radio. Then maybe these lessons learned can be applied to the 60 GHz band where we have the transmit power to do something useful.

RB: Now that you're back at Stanford, what are your

future research directions there?

TM: After some hard times trying to figure out what would be interesting after wireless, I picked a research area in neural signal processing. I am working with Professor Krishna Shenoy, who has a lab at Stanford that has the facility to directly detect individual neural signals using hundreds of electrodes implanted in a primate's brain. We are trying to deduce how information is encoded in the neural signals to better understand how brains function.

RB: Are you understanding what kind of signal processing the brain is doing?

TM: That's what we intend to do. For example, by tapping into a few neurons in the motor-cortex area of the brain and by training the primate to move its arm in a predetermined direction, we can train the neurons and learn from their signals as to what was intended by the brain—which direction, speed, and movement the hand was told to do.

The immediate application is for prosthetic patients

who have suffered spinal-cord injuries. We can detect the brain signal directly and transmit its intention to a robot arm or exercise a muscle of the arm so that the patient can function as normal.

Since my field is in signal processing, my job is to look at a signal, correlate it with its intended goal, and figure out a model that we can use to decode the information embedded in the signal. We want to decode information from neural signals and implement the procedure in a brain implant. The research area that I'm personally involved in is not only on the signal processing side, but also on the circuit side—the circuit technology for a brain implant. What kind of circuit platform, for example, do we need to develop in order to process signals gathered by hundreds or thousands of electrodes in a very small area—highly sensitive but consuming extremely low power—and decode those signals, thousands of them, then wirelessly transmit the decoded information from inside the brain tissue directly to the external world?

RB: Do you think these new algorithms will be useful for future wireless systems?

TM: Oh yes. For example, we talk about cognitive radios, where the radio will have to learn from its environment which spectrum is available and how much power is needed in order to communicate certain pieces of information from one user to the other.

I think this type of computation is not suitable for implementation using a microprocessor. It cannot be programmed effectively, so another branch of the research I'm working on is to learn if we can borrow some of the good ideas that the brain displays in our experiments. We call it plastic computing: how to adapt gradually, but learn from its environment and change its computation fabric automatically—not through software programming but through hardware reconfigurability. I think there's a great opportunity to apply this technology to wireless communication in the future.

By the way, the brain is a very, very dedicated machine. It does not use programs. Each neuron is trained to perform a very specific function. That's how the brain works. Interestingly enough, this is also how silicon works. We have billions of transistors interconnected together, and there is really no need for a central control. Actually, having a central controller usually makes the situation very messy, and that's exactly what we witness in microprocessor design today. You have a CPU and you have to deal with the global clock, and that makes the design overly difficult.

RB: How would you advise someone who wants to get involved with digital signal processing and new applications?

TM: You will have to master a few areas. You cannot just learn math and ignore the silicon, especially in signal processing where the algorithms really mean nothing unless you have an interesting, efficient way of implementing them. Then you immerse yourself in the application domain of your interest.

RB: Do you have any ideas on what you might do after the neural area?

TM: Probably something bio-related. I do feel that signal processing is the basic tool that can be applied to many different areas. We have applied it to low-power circuit designs, video processing, and, most recently, wireless communication. I think in the next several decades signal processing will be widely used in the bio field—for example, genome analysis or diagnostics. Signal processing is after all a science for optimal detection. I think

there might be some interesting developments in those areas.

RB: What causes DSP to move into these different areas?

TM: In the early '80s, all of us learned about DSP processors. The DSP processors solved some of the very simple problems, primarily in the audio band where they could meet the performance and power requirements. But when we moved into the '90s, where multimedia became a big expectation, the amount of processing required far exceeded what the DSP processors could deliver. That's when a few companies and some of the university research programs started to take a look at just exactly what kind of raw computation power CMOS technology could deliver.

We learned to retain flexibility but without programmability. That's when we started to see a lot of interesting products based on DSP: DSL, satellite TV, and then cable modems, set-top boxes, 3D games, high-definition TV, and most recently wireless LAN. I think this trend will continue because digital signal processing matches the CMOS technology very well. It matches the technology better than microprocessors or analog signal processing. Because CMOS circuits are made of fast switches, it's precisely the property that signal processing requires and naturally gives it the kind of unique position to benefit from the technology advancements.

RB: People have been talking about getting near the end of the CMOS roadmap. Do you see that somehow making a big change in things you might work on and what you might do?

TM: No. I think people have doubts about the CMOS roadmap because the architecture that we use today is no longer appropriate.

The von Neumann machine, invented some 60 years ago, was based on certain premises and assumptions that are no longer valid. If we insist on building all computations using centrally controlled processing units, then CMOS probably cannot give us more benefit or performance gains in the future, but if we again let go of the traditional wisdom and take a fresh look at the future of the technology, then I think the computation model and the architecture will be changed and new applications and implementation methods will emerge. ☺

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

© 2004 ACM 1542-7730/04/0300 \$5.00