

Size does matter – the case for big motes

Robert M. Newman, Elena Gaura
Coventry University, Priory St, Coventry UK, CV1 5FB

ABSTRACT

A study conducted by the authors on some contemporary designs for wireless sensor nodes, or motes shows a large disparity in the level of resources, in terms of processor power and available memory. There are essentially two schools of thought on the appropriate scale of hardware to include in a mote design: the 'small mote' camp motivates their approach based on an argument centring on cost, power consumption and potential for integration into a single chip design. The 'large mote' camp cites an expected level of needed computational power, derived from top-down design considerations, as a reason for the scale of the resources designed into their motes. The authors have recently begun to plan some real world, large, field sensing applications, typical of the type of task that may be required for any sensor networks deployed to deliver field maps over real surfaces. In this paper the relative merits of 'big' and 'small' motes are compared in three categories, which are considered by most researchers to be important for the construction of successful wireless sensor networks. The categories are: size, power consumption and cost. In none of these is the big mote found to be at a disadvantage compared to the small mote, and there appear to be definite advantages of big motes with respect to cost and power consumption, at least for the types of application for which large sensor networks may be useful. As a result of these considerations, a specification for a 'big mote' is proposed.

Keywords: wireless sensor networks, motes, system design.

1 INTRODUCTION

As research into wireless sensor networks progresses, there remains a question still widely unanswered – what or where is the 'killer application' that will justify the research? At present, real wireless sensor systems are relatively small scale, and usually employed for scientific investigation, in circumstances in which it is impractical to wire sensors. While undoubtedly useful, these are unlikely to have the impact required for a killer application. The authors believe that the killer applications are to be found in the domain of field sensing, that is the sensing and mapping of physical phenomena in large two or three-dimensional spaces. In these applications a large number of sensors are deployed in an array. It is the ability of such an array to satisfy a number of hitherto unrealizable dreams for the researchers concerned which provides a new potency for environmental, pollution and scientific monitoring systems.

- A sensor can be placed close to some sensed event, wherever it chances to occur
- Sensing or monitoring can go on 24 hours a day, seven days a week.
- The sensing systems can be physically sufficiently unobtrusive that they do not interfere with the daytoday life in the are being sensed.

These systems will not become fully specified until the domain experts in the field begin to appreciate the characteristics of wireless sensor network technology, and lay down the basic functions required for the system. Such specifications are now beginning to emerge, and it is clear that the requirements are quite different from those that have driven WSN research to date, and consequently, the optimum node or 'mote' for such systems will be different from those that have been the mainstay of the field until now.

2 THE BEGINNING: MOTE DESIGN CONSIDERATIONS

Much of the current research on topics related to Wireless Sensor Networks (WSN) pursues an agenda which was set by Kris Pister in 1998: Smart Dust [1]. Smart Dust was a vision, but non-specific about applications. Instead it was motivated in terms of the general promise of ubiquitous computing. Pister relates the vision as follows [2]: "In 2010 MEMS sensors will be everywhere, and sensing virtually everything. Scavenging power from sunlight, vibration, thermal gradients, and background RF, sensors motes will be immortal, completely self contained, single chip computers with sensing, communication, and power supply built in. Entirely solid state, and with no natural decay processes, they may well survive the human race. Descendants of dolphins may mine them from arctic ice and marvel at the extinct technology."

The major constraint in envisaged in the design of these motes was a need for tiny size. As stated in a subsequent paper [3]:

"Size reduction is paramount, to make the nodes as inexpensive and easy-to-deploy as possible. The research team is confident that they can incorporate the requisite sensing, communication, and computing hardware, along with a power supply, in a volume no more than a few cubic millimeters, while still achieving impressive performance in terms of sensor functionality and communications capability."

Along with the tiny size as design goal, the Smart Dust vision introduced a number of other assumptions, mostly as consequences of the view of how WSNs would be deployed, which were in turn derived from studies of

battlefield sensing. They are as follows.

Size

The method of deployment most often mentioned with respect to 'smart dust' motes is by ejecting them from the back of an aircraft. This seems to make sense in the context of battlefield sensing. This mode of deployment calls for a very small mote, since large ones will achieve a high terminal velocity and most likely smash when they hit the ground. Large motes would need some kind of parachute, which in turn would need some mechanism for jettisoning it before the mote landed, in order to prevent the parachutes from interfering with the motes' sensing duties. For this reason, there has been a strong driver in WSN research towards the production of tiny motes.

Power

Along with a very small mote goes very small power resources. Since battery capacity is broadly proportional to volume, the power storage of very tiny motes is likely to be minimal. Also, existing means of gathering power from the environment (for instance, solar cells, RF power harvesting) produce an amount of power proportional to their area, and therefore also do not scale down well. For any extended lifetime it is necessary to either make the mote a nano-power device or to find some means of harvesting power from the environment in the form of vibration energy, or both. Both issues have been important in WSN research.

Cost

The vision that motes would be used in vast quantities, on a 'throw away' basis has led to an assumption that costs must be truly tiny to make such systems practicable. It would appear that the only way to ensure such a tiny cost would be to achieve single chip designs. Hence, the single-chip mote has been a research goal, recently achieved [4]. Often, with very low cost single chip devices, the cost of the chip is exceeded by the cost of its packaging, but not much attention seems to have been paid to this issue, perhaps because no single-chip mote has been contemplated for a real world application.

Although the early drivers for WSN did not, perhaps, include provision for tightly specified, achievable and useful applications, the technology that was produced by those early days has indeed been found to be a useful tool in the real world. As a result, a number of real applications have been deployed, but none match the scale or impact of the scenarios put forward in the early research.

3 A KILOMOTE PERSPECTIVE

As stated in Section 1, the technology is now sufficiently mature to enable very large, field sensing networks, with thousands or tens of thousands of nodes – kilomote systems, and genuine applications using these systems are being scoped. With a better view of real applications for these kilomote systems, it is possible to take a second view of the aforementioned drivers of mote design.

Size

Although it would seem that there is advantage to be had

if motes are small, dust scale motes are a non-starter. In real applications, motes will not be deployed from the back of an aircraft. In any case, in almost any application (except battlefield sensing) it will be necessary to pick up all the motes so deployed, unless they can be made biodegradable, and this has never been a research goal in WSN. A mote must be large enough for its operation not to be prejudiced by environmental hazards such as, for instance, blades of grass. In many applications they will be mounted on (rather large) posts, poles, buoys etc. Generally, these systems will be carefully installed by trained technicians. The real constraint on their size is, as a lower bound, that they must be large enough to be readily handled and reasonably robust. As an upper bound, they must be small enough not to interfere with the visual or physical environment in the space in which they are installed. As a rule of thumb, matchbox size seems to be fine.

Power

The major power constraint is that a mote must be capable of operating for its service life without requiring new sources of power. If service life is taken to be, say, ten years, which may be the requirement for permanent environmental sensing systems, then there is no current battery technology which can provide for this kind of life-span, yet alone the capacity to power a mote for this duration. This means that, in such applications, motes must draw power from the environment. Whether this is a problem is far from clear. In one recent meeting with application specialists the authors were informed that 'there is almost nowhere in the UK where electrical power is not readily available' – clearly not true as an absolute statement, but one that was accurate in the operational experience of the engineer making it. On the other hand, some kind of power is likely to be available in most locations, and its provision is more tractable at matchbox size scales. For example, the best solar cells produce a power output of around 200Wm^{-2} for an irradiance of 1000Wm^{-2} [5] (obviously, 20% efficiency). For a matchbox panel size cell ($2.5 \times 10^{-3}\text{m}$) and a solar irradiance typical of northern Europe in winter [6] (40Wm^{-2}) the amount of power available (continuously) would be 20mW. This provides a ball park estimate for the kind of power consumption that would be reasonable for a real-world mote.

Cost

Installation of a kilomote system, by trained technicians, will not be a low cost undertaking. Even if the motes were without cost, these will be expensive systems. This does not make them impractical, it merely means that for real systems to be deployed there must be genuine returns from that deployment which exceed the cost of these systems. After discussion with some application domain experts, the authors are in no doubt that applications with genuinely deep impact exist, with a sufficiently large return to require them to be kept confidential at present. Thus, it would not

seem necessary to aim for a very low cost mote. For many of the target applications, a cost of around \$100/unit in

The Intel and Sun motes seem to have been designed from a different perspective. They may be termed ‘big

Table 1: Some contemporary intelligent sensor designs

	MICA2	Telos	BSN	EYES	Intel	Sunspot
Size	32x57	32x65	30x30	28x75	30x30	35x40
CPU	Atmel Mega 128	TI MSP430	TI MSP430	TI MSP430	ARM 7	ARM 7 75MHz
MIPS	8	8	8	8	10	70
Program memory	128k	48k	64k	48k	512k	2M
Data memory	4k	10k	256	10k	64k	256k
Logging memory	512k	512k	0	1M	0	0
AD Channels	8	8	6	8	0	0
AD bits	10	12	12	12	n/a	n/a
Wireless			802.15.4 (zigbee)		Bluetooth 1.1	802.15.4 (zigbee)
throughput	38.4kbps	250kbps	250kbps		723.1 kbps	250kbps
OS	TinyOS	TinyOS	TinyOS	AmbientRT	TinyOS	Java VM

1000 off quantities seems reasonable. This does not require an advance on current mote designs, which will be discussed in the next section, although some production engineering would be required to make most of the current designs producible in quantity at this level of cost.

4 CURRENT WIRELESS SENSOR NODE DESIGNS

Although the single chip mote originally envisaged in ‘Smart Dust’ has now been achieved, the majority of useful mote designs today are based on small printed circuit cards, and use COTS technology. The most ubiquitous wireless intelligent sensor design to date is the Berkeley ‘mote’ design, which was originally published in 1998 [7] and has since been continuously refined, both in a commercial direction, and towards a single chip architecture. Table 1, produced from information from Falchi [8], Kling [9] and Crossbow [10] shows the resources available to several generations of mote. The first two are the original Berkeley designs. The BSN mote [11] has been designed by researchers at Imperial College, London, specifically for body sensing applications.

The Intel and Sunspot motes are ‘tasters’ for these two large hardware companies for this new technology.

As stated in section 1, the drivers for the Berkeley design has been the ‘smart dust’ vision of future ubiquitous sensing and a plan of working towards this using available technologies. Currently, most of the feasibility work being done on ‘mote’ designs uses commercial, COTS devices, as do all of the motes in table above.

An inspection of the table shows that mote designers have generally followed two paths, at least after the now aging MICA design. One design is what might be called the ‘small mote’. Broadly, these follow the Berkeley Telos design, which was based on the TI MSP430 16 bit microcontroller – selected primarily on the basis of its very low power consumption. The issue of power consumption is discussed in section 6.

motes’, with 32 bit processors and substantial memory. The motivations of the designers in making this departure are not clear, at least in the papers that have been published (as commercial developments, there has been less disclosed than would have been the case if they had been developed in public research institutions. In fact, the motivation of the designers of the Intel mote seems quite contradictory. While it is a ‘big mote’, it uses the aptly named TinyOS operating system, whereas the SunSPOT mote, even better endowed in terms of processing resources, is limited to a single programming language, Java (probably a consequence of Sun’s commitment to this language). The scale of the Sun mote’s resources have been determined by the requirements of the Java Virtual Machine.

The use of a single language system seems reasonable for the current small mote systems, but are likely to be inadequate for real, kilomote systems. These are likely to include large, sophisticated software systems embedded in the network. Some of these software systems will need to make use of ‘legacy’ codes, developed by applications specialists, and implemented in their preferred language. In the same way that the first generation of single language networked personal computers was succeeded by multiple language ones, single language motes will be replaced by ones which can support a multiple language, general purpose, operating system. Compatibility with legacy TinyOS software can readily be provided by a suitable API layer over the OS kernel. Modern modular operating systems such as eCOS [12] allow multiple APIs, and would appear to be a natural choice. However, it does not seem reasonable, in the long run, for application experts to program bespoke systems in order to install a wireless sensor network. For this reason, there is much current work on higher level means of obtaining information from WSNs, and, if successful, these should provide a potent tool to reduce the cost of ownership of a WSN. Moreover, they will allow the use of the network, once installed, for the provision of information of types not envisaged when it was

designed. Such a need would be typical of WSNs used for scientific exploration, in which the nature of the information required is likely to change as the research proceeds. The authors experience with the development of a query system, called ASQue shows that such a system will not be simple or lightweight if it is to have sufficient power, expressivity and flexibility to be usable over a number of domains. It is already clear that such a system will put substantial demands on the processing and memory resources of a mote.

Drawing on the discussion above, an outline specification suitable for a mote for the new kilomote networks is given below:

1. The mote should support a general purpose, multiple language operating system, preferably open source.

There are many advantages of 32 bit processors with substantial memory resources, as will be discussed in section 6. The 32 bit processor of choice is the ARM, since it is available in a variety of configurations from different vendors (for COTS motes) and is also available as IPR (for ASIC, single chip motes). Moreover, ARM development continues apace, and it is unlikely to become obsolescent in the foreseeable future. Therefore:

2. The mote should utilize the ARM processor architecture, in order to provide a seamless progression from current COTS motes to future ASIC based motes.

The expected type of application, and the requirement to use legacy code and support a multiple language system suggests that the mote should have substantial memory resources. However external memory chips do use additional power, so it is desirable for the memory to be integrated onto the processor chip.

3. The mote's memory (Flash, RAM) should be as large as is feasible, given constraints of availability or chip size.

Currently, the largest memory available integrated with an ARM processor is on the Atmel AT91FR40162S (as used in the Sun SPOT mote), which has 2 Mbyte Flash and 256kByte SRAM.

A continuing emphasis in mote design is power reduction. However, power consumption must be reduced by orders of magnitude before it becomes feasible for a wireless mote to operate for several years from a battery of feasible size. In any case, there are few, if any, battery technologies with a lifetime of years, whatever the power drain. A more sensible emphasis for practical motes, is on available sources of ambient power, to keep cells charged. Therefore:

4. The mote should be equipped with a flexible power conversion circuit, allowing it to be powered from a range of ambient sources. The power management system should allow for the charging of on-board cells for use when ambient power is not available.

Putting these together, it would appear that this mote will be substantially larger, more expensive and power hungry

than existing motes, since it provides so much more processing power and memory resources. It will be argued in the next sections that this need not necessarily be the case. Sections 5, 6 and 7 look separately at issues of size, power consumption and cost.

5 SIZE

The physical size of COTS motes is not primarily determined by the size of the processor chip. Rather, it is constrained by the size of the chip package, and the number of connections made to that package. It would seem natural that a 32 bit processor would require a greater number of connections, due to the wider bus. In fact, all of the motes in Table 1 are designed around systems on a chip, with RAM and flash ROM integrated on the chip. Since there are no external memory connections, the size of the bus is immaterial. The package of the 16 bit, 62 kbyte MSP430 processor used for the Telos mote measures 9mm square, whereas that for the 32Bit 2 Mbyte Atmel ARM measures 10mm square. This it would appear that big motes require little extra PCB space, when compared to little motes. Often the 'small mote' design philosophy is motivated by the need to keep things simple in order to ultimately achieve single chip motes. The first implementation of such a mote is the SPEC mote, a little mote, which is 5mm square [4]. Taking 2004 technologies as a baseline, we can tabulate the chip real estate that will be required for the major subsystems for a 'big mote', assuming a requirement for 1 Mbyte memory.

2.4 GHz radio transceiver	0.25u BiCMOS	8mm ² [13]
ARM 7EJS processor	0.18u CMOS	0.95mm ² [14]
Intel Static RAM cell	0.13u CMOS	5×10 ⁻⁶ mm ² = 40mm ² [15]

Adding these together we come to a total chip area of around 45mm², or 7mm square. This is an easily feasible size for a commodity chip, and could include die space for MEMS components, if this was the integration route chosen, without the chip reaching an uneconomic size. Thus a single chip cogent MEMS sensor is easily possible using established technology.

One interesting point about this is that the processor is the smallest element here, even though this processor, a 200 MIP 32 bit RISC is far more powerful than those commonly found in current sensor motes. On the other hand, memory consumes a dominant share of the chip area. A powerful processor does not consume a substantial amount of silicon real-estate, and could be a good investment, if the additional processing power can reduce the requirement for other specialist hardware. For instance, if the enhanced processor provides for additional digital signal processing, which in turn provides some type of data reduction, then this can save both memory space and communications bandwidth, both of which are critical resources in terms of the wider system design. In designing intelligent sensor devices, it is best to work from a systems point of view, and judge the design options on their merits.

6 POWER

Does a 'small mote' use less power than a 'big mote'? The answer to this question is not simple. In fact, it would appear to depend on the way that the mote is used. Power consumption of a mote arises from the power needs of the data acquisition, on the one hand and those of the RF communications on the other. The usual wisdom is that the requirements of the RF dominate those of process, and this would appear to be borne out by studies such as that by Shnayder et. al [16] which reported on power consumption tests on a Berkeley MICA mote. In these measurements, the active processor consumed 8mA (24mW), while the radio receiver alone consumed 7mA (21mW). Transmitter power was dependent, of course, on the transmit power setting. At maximum power (which, in the authors' experience of these motes, is usually necessary in real networks) the transmitter consumed 21.5mA (65mW). Thus, if both radio and processor are constantly active, the radio uses 86mW, nearly four times the processor's power (and also four times the power available from the solar cell, discussed in section 3). In the quest to save power, the processor and transmitter will be kept inactive, except when needed. When quiescent, the processor consumes 216 μ A (650 μ W). Thus, in a standby state, the mote uses 32 times as much power for communication (the receiver's 21mW) as it does for processing. The communications cost of the transmitter is 27 μ J per byte transmitted. This is an 8 MIPS processor, so at an active power consumption of 24mW, each instruction costs 3nJ. A byte of data transmitted consumes the same power as 9000 instruction executions. Consequently, although processor power is an important consideration, in the overall picture it is a secondary consideration when set against the needs of the RF. Even large reductions in processor power consumption will not have a concomitant effect on the mote's power usage.

The picture is made more complex, since modern embedded processors have power saving modes and variable clock rates. If an efficient power management strategy is used, the processor's power consumption may be made to an extent proportional to the amount of processing actually undertaken. Thus, if the power of the processor is not actually used, it is not paid for.

Power consumption of processors is generally measured in MIPS/W. The MSP430 processor used in the newer Berkeley derived motes (Telos, BSN, EYES) has a particularly good MIPS/W rating of 1333 [17], being specifically designed to be exceptionally power efficient. Of the large processors, the ARM cores are generally believed to be amongst the most power efficient. A recent product announcement by Atmel Corp for the AT91SAM9261 controller claims a MIPS/W rating of 1058[18], lower than, but still in the same league as, the MSP430. The recent clockless ARM996HS core announced by Handshake Solutions [19] is expected to be three times as power efficient as this, and will outperform the MSP430.

However, not all MIPS are the same. The ARM is a 32 bit processor, with many DSP-like instructions, and can perform many tasks using far fewer instructions, and hence power, than can an 8 or 16 bit processor. In many cases the energy saving due to the use of an adequately resourced processor can be enormous. In a design study by Mitra et al [20], the energy consumed by various processors to undertake a 128 point FFT calculation is measured. It is found that the 8 bit processor in the MICA mote requires 930 μ J to perform this task, whereas an Intel PXA255 (an ARM derived processor) requires just 45.8 μ J, about 20 times less energy for the same task. This indicates that 32 bit architectures, such as the ARM may be much more energy efficient for a certain type of task. The reason for the huge inefficiency of the small processor in these tasks is that the range of the numbers involved exceeds those given by 8 or 16 bit registers, so the smaller processors are required to accumulate results resulting in additional instructions and memory transactions. Moreover, they tend to have extremely limited memory and processor registers, so the algorithms employed become very convoluted in order to fit within that which is available. By contrast, the ARM is performing 32 bit multiplications and adds in a single clock cycle, and has sufficient memory to hold the entire FFT frame in one go.

The question is: is the type of task represented by the FFT one which will be typical in large wireless sensor networks? If it is, then the 'big mote' is likely to be a better choice, in terms of energy consumption. The question posed is however very difficult to answer generically without real experience of the applications for which large wireless sensor networks will be used. Current real networks are small, and are designed to be just those functions which are feasible with current (predominantly 'small') motes. However, it is possible to conceive of situations in which the FFT calculation used by Mitra et. al. would be the kind of task operated by a mote in a large WSN. Consider a sensor array deployed in a vibration monitoring system. The task on each sensor is to look out for vibration at characteristic frequencies, and should they exceed some limit, notify the user. This could be done using a set of digital filters, or by performing an FFT. If the diagnosis task is centralized, the mote needs to send 128 bytes of data (assuming a 128 point FFT at 1 byte/sample) for each diagnosis cycle. This will cost 3.5mJ, assuming the use of the Mica RF system. To perform the same diagnosis locally will cost the Mica nearly 1mJ, and a 'big mote' (as represented by the Intel processor in Mitra's study) 50 μ J (the cost of two bytes' worth of communication power). So long as there is nothing to report, no RF power is used. It can be concluded that if this type of task is representative of future kilomote applications, the big mote will consume much less power than the small mote.

7 COST

It is very difficult to make a confident assertion that an 8

or 16 bit processor is lower in cost than a 32 bit one. For instance, the current price of the MSP430 used in the Telos style motes is currently \$5.60 in thousand-off quantities. An ARM processor to a similar specification (memory and peripherals) can be purchased in similar quantities for less than four dollars. The low price of the ARM processor is due to the fact that they are produced in massive quantities – the ARM architecture accounts for the majority of the 32 bit computers in the world today, and the parts are made for high volume applications such as mobile phones, PDAs and network switches. It is likely that the ‘piece part’ cost of a ‘big mote’ is of the same order as that of a ‘little mote’.

8 CONCLUSION

Although no kilomote applications have yet been fully specified, they are now becoming sufficiently visible to infer their general characteristics, allowing the production of a (very rough) outline specification for a mote that can be used in this type of system. This exercise leads to a mote specification which is definitively in the ‘big mote’ camp, with a 32 bit processor and large memory. It has been shown that a ‘big mote’ does not necessarily involve a larger physical size, higher cost or heavier power consumption than the currently preferred small motes. Indeed, the power consumption, of a ‘big mote’ may, in some possibly realistic circumstances be significantly better than a small mote. It is therefore proposed that the ‘big mote’ provides a better component for the building of kilomote systems than does the ‘small mote’.

REFERENCES

-
- [1] K. S. J. Pister, J. M. Kahn and B. E. Boser, "Smart Dust: Wireless Networks of Millimeter-Scale Sensor Nodes", Highlight Article, 1999 Electronics Research Laboratory Research Summary.
- [2] Pister, K., My view of sensor networks in 2010, <http://robotics.eecs.berkeley.edu/~pister/SmartDust/in2010>
- [3] J. M. Kahn and R. H. Katz and K. S. J. Pister , Next Century Challenges: Mobile Networking for "Smart Dust", International Conference on Mobile Computing and Networking ({MOBICOM}), 1999, pp. 271-278
- [4] Pam Frost Gorder, "News," Computing in Science and Engineering, vol. 05, no. 6, pp. 6-9, 2003.
- [5] Sunpower corp, product brochure, <http://www.sunpowercorp.com/html/Products/Solar/solarcells.html>
- [6] World Energy Council, <http://www.worldenergy.org/wec-geis/publications/reports/ser/solar/solar.asp>
- [7] V. Hsu, J. M. Kahn, and K. S. J. Pister, "Wireless Communications for Smart Dust", Electronics Research Laboratory Technical Memorandum Number M98/2, February, 1998
- [8] Anastasi, G., Falchi, A, Passarella, A, Conti, M, Gregori, E, Performance measurements of motes sensor networks, Proceedings of the 7th ACM Int. Symp. on Modeling, analysis and simulation of wireless and mobile systems, 2004
- [9] Kling, R., Intel Mote, Pres. at Berkeley Wireless Research Center, <http://bwrc.eecs.berkeley.edu/Seminars/King-6.27.03/Intel%20Seminar.ppt>.
- [10] Crossbow technology Inc. <http://www.xbow.com/Products/productsdetails.aspx?sid=126>
- [11] <http://ubimon.doc.ic.ac.uk/bsn/index.php?m=185>
- [12] Red Hat, Inc., Garnett N., Larmour, J., Lunn, A., Thomas, G., Veer B., eCOS Reference Manual, <http://ecos.sourceware.org/docs-2.0/>
- [13] Rigge, L and Grewe, T (2003), Tradeoffs Shape Multimode WLAN Silicon Designs, <http://www.commsdesign.com/showArticle.jhtml?articleID=16500948>, Agere Systems, Apr 08, 2003 (6:06 AM)
- [14] ARM Ltd, ARM Technical Support FAQs, <http://www.arm.com/support/faqip/3735.html>
- [15] Yang S, Bai P, Bramblett T, Crew B, Hussein M, Jacob P, Kenyon C, McIntyre B, Moon P, Sivakumar S, Tufts B, Thompson, S., Tyagi, S., and Bohr, M. (2000) A 130nm Generation Logic Technology Featuring 70nm Transistors, Dual VT Transistors and 6 layers of Cu Interconnects International Electron Device Meeting, San Francisco.
- [16] Victor Shnayder, Mark Hempstead, Bor-rong Chen, Geoff Werner Allen, Matt Welsh, Systems Simulating the Power Consumption of Large Scale Sensor Network Applications proc. Conference On Embedded Networked Sensor pp 188 – 200, 2004.
- [17] Texas Instruments Inc, product information, http://focus.ti.com/mcu/docs/overview.tsp?familyId=342&templateId=5246&navigationId=11466&path=templatedata/cm/mcuovw/data/msp430_ovw&DCMP=MCU_other&HQS=Other+IL+msp430
- [18] Atmel Inc, News release, Atmel's New 200 MIPS ARM9 MCU Draws Only 2.5 uA Standby, and 350 uA/MHz at Maximum Performance, http://www.atmel.com/dyn/products/other_docs.asp?family_id=605
- [19] Handshake Solutions, ARM996HS White Paper
- [20] A. Mitra, A. Banerjee, W. Najjar, D. Zeinalipour-Yazti, V. Kalogeraki and D. Gunopulos, High-Performance Low Power Sensor Platforms Featuring Gigabyte Scale Storage, IEEE/ACM 3rd International Workshop on Measurement, Modelling, and Performance Analysis of Wireless Sensor Networks (SenMetrics 2005), , San Diego, CA.