THE UBIQUITOUS SWARM LAB AT UC BERKELEY –

RESEARCH VISION IN A NUTSHELL

BERNHARD BOSER, BJÖRN HARTMANN, EDWARD LEE, JOHN KUBIATOWICZ,
KEN LUTZ, KRIS PISTER, JAN RABAEY, CLAIRE TOMLIN (AND MANY OTHERS)

JULY 2012

This document presents in a nutshell the ideas that have led to the formation of the Ubiquitous Swarm Lab at Berkeley. They are the results from discussions among the Center members as well as interactions with many leading researchers in the field of ubiquitous (pervasive) distributed computing, communication, and control. It builds on the decade-long experiences of the Center members in the fields of wireless sensor networks and cyberphysical systems [REF].

The "Swarm" finds its origin in the *wireless sensor network* (WSN) concept that emerged in the mid-1990s [Pister]. Enabled by the introduction of low-cost and low-power technologies for sensing, data acquisition, computation, and communication as well as the emergence of ever more pervasive wireless networking, WSNs were perceived to have the potential to dramatically impact the way large societal systems were implemented and operated. Environmental and infrastructural monitoring, building and traffic management, power distribution, health care, industrial automation, and emergency response are just a few of the many areas that could be transformed by the capability of having continuous and distributed monitoring combined with data aggregation and closed-loop control. The academic community embraced the WSN concept with enthusiasm as it opened a Pandora's box of interesting and challenging new problems. The business community, on the other hand, reacted with far more reserve. While indeed there are a number of high-visibility deployments in each of the aforementioned domains, business volume is small and represents only a tiny fraction of the overall information-technology revenue. Yet, after some hesitance about scale and feasibility, most companies in the related business areas have now started to embrace the idea.

The obvious question at this point is what has hampered a widespread adoption of the concept so far, and if and how these hurdles can be overcome. A number of technological issues can be readily identified:

- Reliability – The applications in question often require 24/7 no-fault operation because of criticality or cost considerations. Ensuring this in a distributed wireless environment under all conditions is absolutely non-trivial.
- Lifetime – Nodes need to operate for a long time without battery replacement or energy replenishments. Accomplishing this in a broad setting may still require one or more orders of power reduction.
- Ease-of-use – The deployment, operation, and management of a distributed WSN network today requires insight and understanding of all the layers of the implementation stack (from the sensor, computing platform, and network to the application), clearly limiting the technology reach.
- Security/privacy – Living and operating in an environment that is fully instrumented (and even controllable) automatically raises profound questions about security and safety, as well as intrusiveness and breach of privacy.

- Cost or disruptiveness – To appeal, a new technology should be either substantially cheaper than the technology it tries to displace, or enable functionality that was not possible before. Neither of these may have been the case as of today.

All of the above are technology issues that can be overcome by time. A more fundamental hurdle plaguing the WSN concept is its *inherent lack of economy-of-scale* (or the potential thereof). The operational WSN model is to create application-specific solutions, where technologies are developed and integrated in a vertical fashion from the application down over software platforms and networks to individual sensors. Hence the realization of networks for power management, environment control, traffic management, etc. While this approach may possibly have efficiency and reliability advantages, it also leads to stovepiping – that is, components of any of these networks can hardly (if at all) be used for other applications that operate in the same space. This has led to the development of a broad range of components, modules, and platforms that are incompatible and generally hard to reuse.

Consider now an alternative picture. Imagine a world flooded with "cyber-interface" nodes all interconnected in a vast global network. The term "cyber-interface" covers anything from input /sensory to output/actionary devices that interface the physical and biological worlds with the cyber world. In such an environment, any "application" could opportunistically exploit the input/output resources available *at the given time and place* to create the best possible utility[1] or experience. A person entering a room and wishing to interact with information could, for instance, borrow the interactive input/output modalities of the room (microphones, gesture tracking, touch surfaces, speakers, displays) rather than rely on the minimalistic user interfaces of her mobile device. The same holds when moving through the town, either on foot or by vehicle. Applications must be interpreted in the broadest sense of the word here, and could be any of the following of the classes of WSN apps described earlier: a building trying to optimize convenience and/or efficiency, a city managing its energy distribution, traffic flow or air quality, an airport managing its security, or a person managing the complexities of daily life. It is this opportunistic mapping between applications and resources that differentiates the "swarm" concept from the "wireless sensor network" and "internet of things" models. Not only does it offer a broad platform vision that enables resources (sensing/actuation, networking, computing, storage) to be shared among many applications, it also allows for a far more efficient utilization of these resources. In addition, the platform abstraction makes it easier for new applications to be created and deployed—as the intricacies of the hardware platform are abstracted away from the application developer—unleashing their creative powers to the fullest extent.[2]

One of the many possible outcomes enabled by the swarm platform could be the demise of the mobile as we know it. In a world immersed in a sea of rich input/output modalities, there is no need for humans to rely on portable devices to seamlessly interact with information, the environment, or each other. In this "unPad" vision[3], the functionality of the pad will remain, but in the form of unpackaged communication, computation, and storage.

However, developing a truly general and open swarm abstraction is a daunting task, requiring innovations at virtually every layer of the system stack. Relevant applications operate under

---

[1] Utility is an application-specific measure of benefit or value to the end user.

[2] Similar to the way the introduction of abstractions such as DOS, Windows, iOS and Android have opened up the PC and SmartPhone platforms.

[3] More information on the unPad concept is provided in the BWRC white paper "The unPad – A Research Vision for the Berkeley Wireless Research Center," July 2012.

tight constraints in terms of latency, reliability, safety, and security, as stated earlier. These have to be guaranteed in an environment where the available resources such as sensing, connectivity, computing, and storage are heterogeneous, distributed, and sparse. Even more, with the operational environment being very dynamic, the availability of resources may change continuously and is strongly dependent upon the scope of the applications of interest. To address these concerns, we propose the development of a mediation layer called the "swarmOS", whose goal it is to broker, provide, and realize service guarantees to competing applications running on the swarm platform, given the resources at hand at a given time and place. It is our belief that a universal, open, and shared platform such as the swarmOS is exactly what is needed to unleash the full potential of the swarm concept.

Given the many challenges, a concerted large-scale effort is the only way to make it happen. The Ubiquitous Swarm Lab at the University of California is doing just that. Seeded by a generous donation from Qualcomm, Inc., the lab brings together in a unique setting, experts in advanced wireless sensor platforms, software operating systems, services control and signal processing, security, innovative user interfaces and applications. Located on the 4th floor of Cory Hall, the open floor plan of the Lab allows for multidisciplinary collaboration and truly immersive research. The ambitious research agenda, driven by compelling drivers such as the unPad, is funded by a consortium of companies from different business segments complemented by grants from Federal and State government. For more information, please consult the Swarm Lab website at http://swarmlab.eecs.berkeley.edu.