Swarm Technology

Edward A. Lee and Jan Rabaey
EECS Department
University of California at Berkeley
Berkeley, CA 94720-1770

Abstract: *The TerraSwarm Research Center is addressing the huge potential (and associated risks) of pervasive integration of smart, networked sensors and actuators into a connected world. Pervasive connectivity, data aggregation, and integration of sophisticated learning and optimization algorithms enable a new generation of systems and services. This paper describes how this work is central to the development of smarter defense systems.*

Keywords: Internet of things; cloud computing; swarm technologies; sensor networks; distributed computing; security and privacy; big data; model-based design.

Introduction

Over the past two decades, there has been a growing realization that large numbers of sensors dispersed into the environment can help to solve societal and military problems. These sensory swarms (as they were called by the second author in a keynote talk at the Asia and South Pacific Design Automation Conference in 2008) can be wirelessly interconnected and interact with the cyber-cloud, and offer an unprecedented ability to monitor and act on a range of evolving physical quantities. Such pervasive observations and measurements enable unprecedented learning and modeling of the physical world under dynamically changing conditions.

At the core are advances in design and manufacturing technologies, which have enabled a dramatic reduction in cost, size, and power consumption of a variety of sensing and actuation devices, along with the familiar improvements in computation, storage, and wireless communication. Industry observers predict that by 2020 there will be thousands of smart sensing devices per person on the planet (yielding a "tera-swarm"); if so, we will be immersed in a sea of input and output devices that are embedded in the environment around us and on or in our bodies.

The concept of wireless sensor networks is not new. Sensor-based systems have been proposed and deployed for a broad range of monitoring (and even actuation) applications. But the vast majority of those are targeting a single application or function. The potential of swarms goes far beyond what has been accomplished so far. When realized in full, these technologies will seamlessly integrate the "cyber" world (centered today in "the cloud") with our physical/biological world, effectively blurring the gap between the two. We refer to such networked sensors and actuators as the "swarm at the edge of the cloud,"[1] and the emerging global cyber-physical network as the "TerraSwarm," encompassing trillions of sensors and actuators deployed across the earth.

TerraSwarm applications, which we call "swarmlets," are characterized by their ability to dynamically recruit resources such as sensors, communication networks, computation, and information from the cloud; to aggregate and use that information to make or aid decisions; and then to dynamically recruit actuation resources — mediating their response by policy, security, and privacy concerns.

Achieving this vision will require a three-level model. The cloud backbone will offer extraordinary computing and networking capability, along with global data analytics, access, and archiving. Mobile battery-powered personal devices with advanced capabilities will connect opportunistically to the cloud and to nearby swarm devices, which will sense and actuate in the physical world.

Ubiquitous connectivity between the cloud and mobile devices such as smartphones is almost a reality today. Through common and general programming and communication interfaces (e.g., "app" programming and TCP/IP) this connectivity has turned the cloud/mobile universe into a flexible platform enabling millions of applications that we could not have imagined a few short years ago. These parts of the system will continue to develop rapidly under large-scale commercial investment. The swarm level, however, because it directly interacts with the physical world, presents challenges that demand forward-looking research. The potential payoff of such research is a system that can fundamentally change and empower human interaction with the world.

Current "smart" applications, such as smart homes, smart grids, and battlefield management systems, typically address a single application on a dedicated set of resources. While this approach provides performance guarantees and reliability, it prevents economies of scale, and, more importantly, it prevents the explosion of possibilities that results from sharing data and devices across applications. The TerraSwarm vision cannot be achieved by a single vendor providing the components as an integrated system.

[1] This phrase was coined by Rabaey in a keynote talk at the VLSI Circuits Symposium in Kyoto, June 15, 2011 [3].

What is needed instead is the swarm equivalent of the common, general, "app" framework that has recently enabled smartphones and similar devices to rapidly deploy and serve a vast range of often unanticipated applications by recruiting resources and composing services. The swarm will never achieve its potential without a "SwarmOS" on which such swarmlets can be built and composed by millions of creative inventors.

While open architectures with dynamically recruitable resources can open up significant security and privacy risks, they can also make systems more efficient (through sharing of resources), more resilient (through dynamic reconfiguration leveraging redundant resources), and more capable, enabling applications we have not yet invented or that cannot yet be realized. Adaptability, reliability, robustness, and security are essential ingredients to be considered from the start.

When the web was first launched, few people would have predicted the astounding range of applications that it would enable. It has profoundly changed the way people interact and behave, how businesses are run, and how information is exchanged. A similar revolution happened with the introduction of mobile platforms such as Android and iOS. We believe that swarm-based systems can have at least as much impact. Enabling this requires a collaborative environment in which to address the TerraSwarm's extraordinarily wide range of challenges and opportunities. By viewing key challenges through many different eyes, we expect to be able to generate a broad range innovative ideas and solutions.

The TerraSwarm Challenge

While the TerraSwarm vision holds enormous promise, it also poses a number of daunting challenges. The technical challenges are defined by the following unique combination of characteristics of TerraSwarm systems:

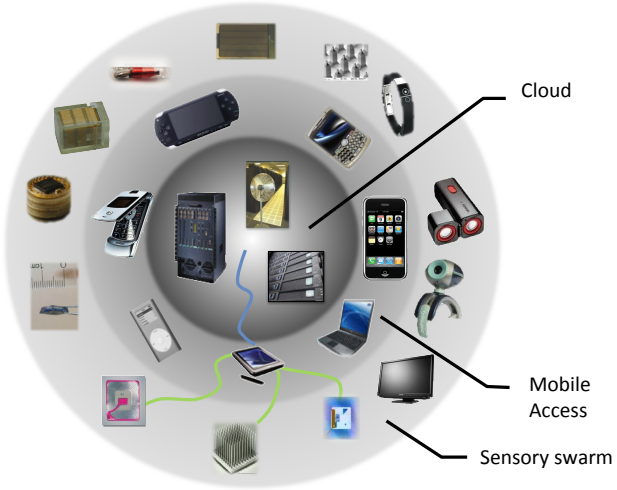

Figure 1. Three-tiered structure of the emerging information technology platform [4].

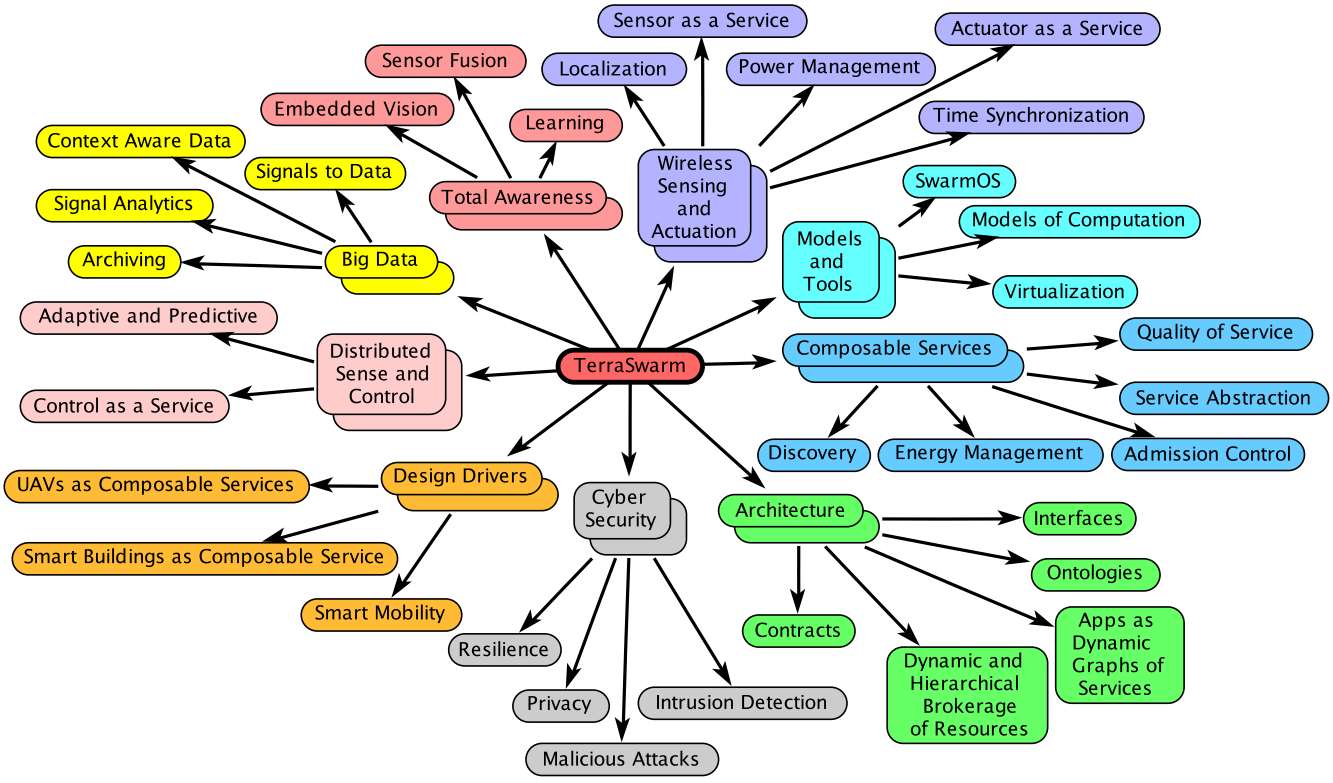

Figure 2. A map of the technical problem space of the TerraSwarm.

- *Large-scale*: the swarm comprises a vast number of nodes generating corresponding "big data;"
- *Distributed*: components of the swarm are networked, separated physically and/or temporally;
- *Cyber-physical*: the swarm fuses computational processes with the physical world;
- *Dynamic*: the environment evolves continually;
- *Adaptive*: the system must adapt to its dynamic environment, and thus the distinction between "design-time" and "run-time" is blurred; and
- *Heterogeneous*: swarm components are of various types, requiring interfacing and interoperability across multiple platforms and models of computation.

Given these characteristics, a sketch of the technical problem space facing a scalable and universal realization of the vision is shown in Figure 2.

The challenges and opportunities include the following:

- Swarm systems rely on vast numbers of heterogeneous sensors that are generating massive amounts of data. How will these data be accessed, processed, stored, and interpreted? First, we observe that data are more valuable when aggregated than when isolated. The emergence of the social networking industry is a case in point. Second, we observe that data need not be communicated or stored if they can be predicted from models. If such models can be learned in an unsupervised way, then the TerraSwarm can be reflective, monitoring its own health, as well as the health of physical devices and humans that it interacts with.
- When data are used for security-or safety-critical systems, how can we verify that they are accurate (i.e., that the sensors are functioning properly), that they have not been compromised (i.e., secure from deliberate or inadvertent tampering), and that their source is known? We observe that today's mechanisms for identity and key management likely will not scale well to the TerraSwarm. The emergence of ubiquitous clock synchronization (with IEEE 1588 and 802.1AS for wired and wireless networks) offers unique new opportunities for scalable security mechanisms, since stable local clocks provide a natural root for trust.
- TerraSwarm applications are generally cyber-physical systems that involve physical actuation and closed-loop control, and hence will have stringent testing and verification requirements. But they will also be highly dynamic, adapting their structure and recruiting resources on the fly. How can testing and verification extend to continuously evolving systems? How can we ensure that effects on the physical world are safe? We observe that on-line verification of adaptive and evolving systems will require lightweight formal methods, something that remains elusive today.
- Swarms and swarmlets that dynamically recruit resources will compete for those resources. How will costs (energy, opportunity cost, and capital investment) be managed? How can we ensure that new deployments do not disrupt established services? We observe that networking innovations such as AVB offer more control over quality of service than has been available in the past on open public networks, but how the control gets exercised in an open and competitive world remains an open question.
- How will we address data privacy and safety? We observe that, counterintuitively, privacy may be easier to preserve with more data than with less, using for example the notion of differential privacy [1].

It is worth observing that nearly every science and engineering university and a broad fraction of industry are engaged in activities that are either directly or peripherally related to swarm systems, often under the heading of the Internet of Things (IoT), the Internet of Everything, Industry 4.0, the Industrial Internet, Smarter Planet, Machine to Machine (M2M), TSensors (Trillion Sensors), or The Fog (like The Cloud, but closer to the ground). Relevant research areas include sensor technologies, actuators, semiconductors, communication systems, control systems, robotics, data analysis, data mining, modeling and simulation, operating systems, energy efficiency technologies, machine learning, data security and encoding, and cyber-physical systems, among others. Thus far, there has not been a coherent effort to bring together these disparate research efforts to serve swarm-based application development. Yet the swarm will only reach its full potential when it becomes a unified, standardized platform enabling the unencumbered development of many swarm applications.

TerraSwarm Research

The nine-university cooperative TerraSwarm Research Center is organized along four themes:

1. Smart Cities. The goal of TerraSwarm is not to develop applications. It is to build infrastructure that enables millions of creative applications. This theme focuses on integrating capabilities developed in the other themes, putting them together to create more sophisticated infrastructure out of the pieces, and to create applications that test and demonstrate the concepts, tools, and software being developed. This is a bottom-up, rather than top-down effort. (For an excellent critique of top-down smart city visions, see [2].) Instead of defining one integrated challenge problem, the goal is to identify and organically grow services and capabilities that contribute to a Smart City, but more importantly, test and demonstrate TerraSwarm technology.

2. Platform Architectures and Operating Systems. The goal of this theme is to develop a hierarchical and compositional system architecture, supported by a distributed, loosely coupled executive that we call the "SwarmOS." This architecture must accommodate heterogeneous and

dynamic compositions of sensor and actuator devices, mobile vehicles, handheld devices, networking components, and cloud infrastructure. A key challenge is to dynamically balance the needs of distributed concurrent application resources, quality of service (QoS), and real-time guarantees. Equally important is the need to respect the privacy and integrity of streams of information. Building on the support of the SwarmOS, TerraSwarm applications will be structured as dynamic, hierarchical graphs of services, with components providing assured quality of service through resource brokerage.

3. Services, Applications, and Cloud Interaction. The goal of this theme is to provide technologies for scalable, adaptive composition of heterogeneous services. Specifically, we envision TerraSwarm applications that combine mobile and fixed sensor and actuator resources that interact directly with physical assets and humans; handheld communication, sensing, and computing devices; wireless and wired networking devices; and networked computing services (e.g. cloud-based computing). A key objective is to enable TerraSwarm applications to leverage large data streams through learning and inference, while preserving privacy and security.

4. Methodologies, Models, and Tools. The complexity of TerraSwarm systems and their safety requirements pose significant challenges for the design of sensing, control, and actuation infrastructures, as well as for application development and deployment. These challenges are compounded by the requirement for on-line adaptation and reconfiguration. TerraSwarm applications need to adapt to the disappearance of resources, recruit useful resources that appear, and adapt services dynamically as part of a utility-driven optimization. There will be less of a distinction between "design time" and "run time," so design techniques, tools, algorithms, and flows must themselves become services that can be recruited online. This is a far bigger challenge than has been previously addressed in the design technology community. One key consequence is that design techniques must be formal and rigorous, or they will not be reliable for on-the-fly reconfiguration, in which there is no opportunity for extensive testing prior to deployment.

Defense Applications

There are many obvious defense applications of TerraSwarm technology, including autonomous vehicle coordination, situation awareness, and advanced weapon systems. The Navy, for example, has a program to develop autonomous boats that coordinate to protect a fleet. Such coordinated autonomous vehicles are central to TerraSwarm research.

But there are also a few less obvious possibilities. For example, classical defense systems use closed, protected networks. But given the capabilities and ubiquity of open public networks, failing to leverage them can only result in tactical disadvantage. And adversaries will use (and have used) open public networks against us (consider the use the cellular telephone network for IEDs). The challenges in DoD use of public networks, however, are daunting, including how to protect information and how to ensure reliable system operation. Preservation of privacy on open, connected networks is all about keeping control over information. And adaptive reconfiguration of swarmlets is all about robustness. We believe that these technologies can enable effective leveraging of open public networks for some defense applications.

Conclusions

Progress towards the TerraSwarm vision requires an astounding breadth of expertise, in large-scale, adaptive, cyber-physical control systems; programming models and tools for heterogeneous, real-time, and distributed cyber-physical systems; security in systems with dynamic topologies; machine learning; privacy; networked sensor and actuator platform design; signal analytics; wireless networking and distributed systems; system architecture; human-computer interaction; energy-aware system design; and application platforms. Only a truly multidisciplinary approach will bring the TerraSwarm vision to reality.

Acknowledgements

The work described here is funded by the TerraSwarm Research Center, one of six centers administered by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

References

1. C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In S. Halevi and T. Rabin, editors, *Theory of Cryptography*, LNCS 387, pages 265–284. Springer, 2006.
2. A. Greenfield. *Against the smart city (The city is here for you to use)*. Do Projects, New York City, 2013.
3. J. M. Rabaey. The swarm at the edge of the cloud - the new face of wireless (keynote presentation). In *Proc. Symp. on VLSI Circuits*, pages 6–8, Kyoto, 2011.
4. J. M. Rabaey, D. Burke, K. Lutz, and J. Wawrzynek. Workloads of the future. *Proceedings of the IEEE*, 25(4):358–365, July-August 2008.