Wireless, Self-Organizing Cyber-Physical Systems

Brian D. Noble and Jason Flinn
University of Michigan, EECS Department

1. Introduction

Cyber-physical systems carry several unique challenges, stemming from their scope and scale. Because they must "go where the work is," they must depend on wireless communication at the edges. There are many devices, and they must move with the physical entities that they measure and control. Furthermore, the real world has physical limitations, and some of its entities will be independent actors with different, possibly competing needs; this is particularly true of the *people* who play a part in the system.

The net result of these two observations is that several components in cyber-physical systems cannot be perfectly controlled. In a wireless network, one cannot prevent nodes from traveling to the edge of connectivity. Environmental considerations—rain and terrain features—cause significant variations in wireless bandwidth. Additionally, one cannot prevent mis-configured or malicious devices from polluting the wireless spectrum. At the same time, the overlay argument [Katz96] suggests that a single, wide-area, high-speed network is unlikely; rather, we will depend on a variety of smaller wireless provisioning points, making movement of devices, along with power management, first-class concerns.

Likewise, the physical entities in our system will not be perfectly controllable. Sensors and actuators suffer inaccuracies due to physical-world limitations, and these physical components can fail. Often, such failures will be partial, providing a failure semantics much more complicated than the familiar fail-stop world of most hardware devices. These devices may also be independent—particularly if the "devices" are the people using the system. For example, an intelligent automotive system is still subject to the will of its human driver. Cultural and legal restrictions will preserve this status quo for the foreseeable future. These users will often have conflicting, competing needs and desires, obviating the possibility of centralized planning and control.

In addition to the challenges of control, the scope and scale of cyber-physical systems gives rise to another significant challenge: respecting the privacy of the people who make use of them. Even now, there are significant exposures of behavior we typically consider private. For example, centralized measurement of a campus-wide wireless deployment makes tracking individuals possible at building and room granularity [Yoon06]. Furthermore, traditional notions of security, authentication, and authorization are unlikely to provide much traction in the cyber-physical world.

2. Adaptive, Proactive Systems

Wireless systems must react to changes in available bandwidth, battery energy, and other unanticipated contextual cues [Flinn99, Noble97]. However, while they appear to be chaotic, there are often underlying regularities in their connectivity [Mickens06] and mobility [Yoon06]. These regularities can be exploited in a variety of ways, improving connectivity [Nicholson06], proactively scheduling opportunistic data exchanges, and provisioning geo-caching, replication and risk–reduction systems.

It can also be beneficial to predict connectivity and contextual changes for application behavior. For example, consider a cooperative, intelligent traffic management systems, using inter-vehicle communication. Control depends on bandwidth, and higher speeds require tighter bounds on round trip control messages. When connectivity conditions are expected to degrade, the vehicles can slow down in *anticipation*, preserving the safety of the transportation system.

3. Self-Organizing Systems

The scope and scale of these systems suggest that a single, planned deployment is simply impossible: devices are too numerous, mobile, and ephemeral for such an approach to have any hope whatsoever. Further, it is unlikely that altruism can be counted upon for correct system behavior. For example, in a loose cooperative of self-interested parties, there is little natural incentive for one to provide resources for another's benefit, and there is strong evidence that *free-riding* occurs in such systems [Adar00]. Finally, while reputation systems may help, they are complicated by the fact that nodes in these systems will likely be able to trivially create new identities for themselves [Douceur02].

Rather than try to plan centrally, hope that altruism prevails, or attempt to audit the behavior of each node, one can instead turn to *incentive-centered design*. In this approach, one builds a system in which a rational node finds it beneficial to make decisions that also favor the system as a whole. For example, cooperative systems can be designed to discover or manufacture pair-wise exchanges, providing fairness without requiring strong identities or centralized control [Cox03]. Similar design techniques can provide leverage in building secure, reliable systems [Wash06].

In addition to self-organization along macroscopic behavior, one can apply techniques to individual nodes to improve the performance, reliability, and efficiency of computing systems. For example, by allowing applications to expose their intent, one can significantly improve power consumption without harming performance [Anand05]. Likewise, one can use such declarations of intent to select network services best matched to application needs [Nicholson06].

4. Secure, Privacy-Respecting Systems

The ubiquity of cyber-physical systems gives them the potential to expose the many details of our lives that we presently consider private. Furthermore, their ability to control aspects of the physical world provides significantly more leverage to a malicious attacker. We must build systems that respect the privacy of their participants without unduly compromising functionality, and re-think our notions of security, authentication, and authorization.

For most people, privacy is not an absolute; rather it is a good that they are willing to trade for other goods and services [Acquisti05]. However, their willingness to expose their own behavior depends on the context in which they find themselves [Ackerman01]. Clearly, it is untenable to ask the user to tell us when his privacy preferences change; instead we must try to infer such changes by tracking context over time. There have been some early successes in doing so [Smith05], but much work remains.

Likewise, traditional notions of authentication and authorization do not translate well to the cyber-physical world. For example, the large number of devices and users means individual device/user authentication is untenable. However, one can exploit location, proximity, and context to be able to automatically manage security contexts [Corner05]. Similarly, the large number of devices necessarily belong to many different administrative domains of control, and establishing security contexts among them is likely to be impossible. Instead, one can leverage everyday activities to help establish secure contexts between collections of correspondents [Nicholson06].

5. References

[Adar00] E. Adar and B. A. Huberman. *Free riding on Gnutella*. First Monday, 5(10), October 2000.

[Acquisti05] A. Acquisti and J. Grossklags. *Privacy and rationality in individual decision making.* IEEE Security and Privacy, 3(1). Jan-Feb 2005.

[Ackerman01] Ackerman, M., Darrell, T., and Weitzner, D.J.: *Privacy in Context.* Human-Computer Interaction 16, 2001

[Anand05] M. Anand, E. B. Nightingale, and J. Flinn. *Self-Tuning Wireless Power Management*. Wireless Networks, 11(4), July 2005.

[Corner05] M. D. Corner and B. D. Noble. *Protecting file systems with transient authentication*. ACM Wireless Networks. 11(1-2), January, 2005

[Cox03] L. P. Cox and B. D. Noble. *Samsara: honor among thieves in peer-to-peer storage.* Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles. October, 2003.

[Douceur02] J. R. Douceur, *The Sybil Attac*k, The First International Workshop on Peer-to-Peer Systems, March, 2002

[Flinn99] J. Flinn and M. Satyanarayanan. *Energy-Aware Adaptation for Mobile Applicaiotns*. Proceedings of the 17th ACM Symposium on Operating Systems Principles, December 1999.

[Katz96] Katz, R. and E. Brewer. *The Case for Wireless Overlay Networks*. in Proceedings 1996 SPIE Conference on Multimedia and Networking. 1996.

[Mickens06] J. W. Mickens and B. D. Noble. *Exploiting Availability Prediction in Distributed Systems.* In Proceedings of the 3rd USENIX/ACM Symposium on Networked Systems Design and Implementation. May 2006.

[Nicholson06] A. J. Nicholson, Y. Chawathe, M. Y. Chen, B. D. Noble, and D. Wetherall. *Improved Access Point Selection*. In Proceedings of the 4th Annual ACM/USENIX Conference on Mobile Systems, Applications, and Services. June 2006.

[Nicholson06] A. J. Nicholson, I. E. Smith, J. Hughes, and B. D. Noble. *LoKey: leveraging the SMS network in decentralized, end-to-end trust establishment.* The 4th International Conference on Pervasive Computing. May, 2006.

[Noble97] B. D. Noble, M. Satyanarayanan, D. Narayanan, J. E. Tilton, J. Flinn, and K. R. Walker. *Agile application-aware adaptation for mobility*. Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles. October, 1997.

[Smith05] I. Smith, S. Consolvo, A. LaMarca, J. Hightower, J. Scott, T. Sohn, J. Hughes, G. Iachello and G. D. Abowd. *Social Disclosure Of Place: From Location Technology to Communication Practices*. In The 3rd International Conference on Pervasive Computing (Pervasive). 2005.

[Wash06] R. Wash and J. MacKie-Mason. *Incentive-Centered Design for Information Security*. 1st USENIX Workshop on Hot Topics in Security. July, 2006.

[Yoon06] J. Yoon, B. D. Noble, M. Liu, and M. Kim. *Building realistic mobility models from coarse-grained traces.* In Proceedings of the 4th Annual ACM/USENIX Conference on Mobile Systems, Applications, and Services. June 2006.

6. Biographical sketches

Brian Noble is an Associate Professor in the Electrical Engineering and Computer Science department at the University of Michigan. His research centers on software supporting mobile devices and distributed systems. He completed the PhD in computer science at Carnegie Mellon University in 1998, and is a recipient of the NSF CAREER award. Email: bnoble@umich.edu. Telephone: +1 (734) 936-2971

Jason Flinn is an Associate Professor of Electrical Engineering and Computer Science at the University of Michigan. He received his Ph.D. from Carnegie Mellon University in 2001, and was granted an NSF CAREER award in 2004. Jason's research interests include mobile computing, operating systems, and storage. Email: jflinn@umich.edu. Telephone: +1 (734) 936-5983