Security for Pervasive Computing

CS239

Kevin Eustice
V. Ramakrishna

4/24/06

What is Pervasive Computing?

One Vision of Pervasive Computing

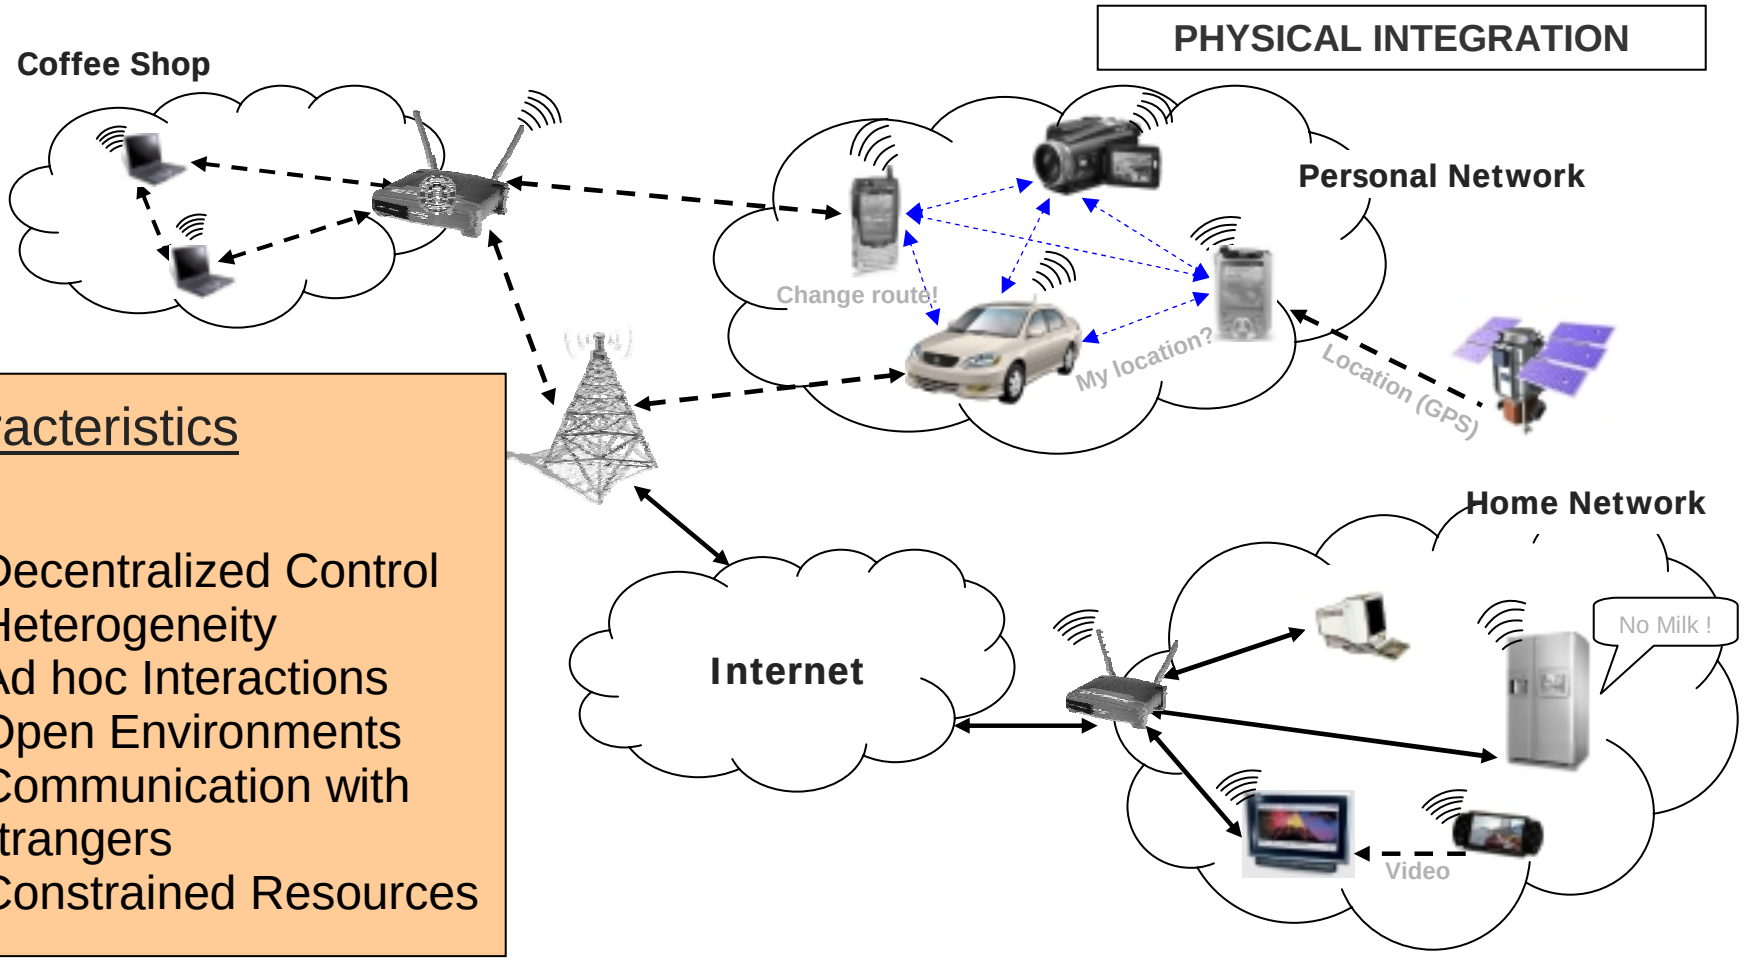

Characteristics

- Decentralized Control
- Heterogeneity
- Ad hoc Interactions
- Open Environments
- Communication with Strangers
- Constrained Resources

SPONTANEOUS INTEROPERATION

New twists on old problems

- Authentication
  - Verifying the identity claims of strangers
- Integrity
  - Protecting mobile devices and data
- Privacy
  - Minimize exposure of sensitive information
- Access Control
  - How do we deal with unknown entities?
- Administrative Challenges
  - Naïve and impatient users

Challenge: Authentication

Problem: *Mutually unknown entities need to verify each other's identity.*

How do we safely add our new wireless device to our home network?

How do we know which device we're interacting with when there are no wires? Is it the AP in the corner? Or is it a PDA in a backpack slung over a chair?

The Resurrecting Duckling

- Addresses the problem of securely adding devices to a ubiquitous computing environment
  - Transiently or permanently
- A computer imprints an identity (a shared secret!) onto a duckling through a *physically secure channel*

Stajano, F. and Anderson, R. “The Resurrecting Duckling: Security Issues for Ad-hoc Wireless Networks” 7th Intl. Workshop on Security Protocols, 1999.

Stajano, F. “The Resurrecting Duckling – What Next?” 8th Intl. Workshop on Security Protocols, 2000.

Talking to Strangers
Bootstrapping Trust Between Strangers

- Extends Duckling model, but removes requirement of physical secure channel
- Exchange *public keys* over a *location limited sideband channel*
  - Unidirectional or bi-directional yields varying degrees of possible authentication
  - “location limited” channel provides some assurance as to actual locality of participants (or their agent)
- What about man-in-the-middle?

Balfanz et al. “Talking to Strangers: Authentication in Ad-Hoc Wireless Networks.” 2002 Network and Distributed Systems Symposium (NDSS 2002).

Balfanz et al. “Network-in-a-Box: How to Set Up a Secure Wireless Network in Under a Minute.” 13th USENIX Security Symposium, 2004.

Challenge: Device Integrity

Problem: How do we protect our devices and networks from malicious users and devices?

Issues:

- Infected or vulnerable devices
- Device or data theft

QED

- Addresses the problem of potentially harmful devices entering your network
- Quarantines devices upon detection, isolating them from other network members
- Examine devices for potential problems – undesirable services/versions, worm signatures, etc.
  - More aggressive internal scans possible w/device cooperation – viruses, packages, etc.
- Decontaminate – repair problems if possible

Eustice et al. “Securing WiFi Nomads: The Case for Quarantine, Examination, and Decontamination” New Security Paradigms Workshop 2003.

Zero-Interaction Authentication

- Addresses problem of device theft or “borrowing” for purposes of data access or modification
- File system contents are kept in an encrypted state
- File access requires user interaction and proximity to a short-range wireless token
  - Decryption information is kept only as long as necessary
- In the case of theft, file system is protected

Corner, M. and Noble, B. “Zero-Interaction Authentication” Mobicom 2002.

Privacy

- Why is this a bigger problem in pervasive computing?
  - We sometimes *have* to interact with strangers
  - Exposure of private information is inevitable
- Wireless Eavesdropping
  - Problems inherent in
    - Broadcast medium
    - Static location of access points
    - Knowledge of 802.11 protocols
- Location Privacy
  - Prevent inadvertent leak of current location
  - Obfuscate location information sent to others

Location Privacy

- *Least privilege* a generally agreed upon paradigm
  - Precision and quality of disclosed data varies, depending on the relationship with the recipient
- Grade information into multiple fidelity levels
  - Arranged in a hierarchy (partial-order)
  - Precision decreases with height
- Access control in Aura[1]
  - Allow users to define policies at informational level
  - Associate per-node access control policies
  - Result: location info from source is obfuscated by the time it is received at the sink
- Negotiate to balance service quality and location privacy[2]

[1] U. Hengartner and P. Steenkiste, "Access Control to Information in Pervasive Computing Environments," *Proc. of 9th Workshop on Hot Topics in Operating Systems (HotOS IX),* Lihue, HI, May 2003, pp. 157-162.

[2] M. Duckham and L. Kulik, "A Formal Model of Obfuscation and Negotiation for Location Privacy ," *Pervasive 2005*, Munich, Germany, May 8-13, 2005.

Access Control in Pervasive Computing

- Issues
  - Scalability
  - Dynamism and flexibility
  - Proof generation
  - Revocation
  - Impact, or side-effects
- Flexible approaches
  - Use roles and policies
  - Separate semantics of roles from their definitions

Generalized Role-Based Access Control

- *Roles* traditionally used for entities
- GRBAC extends the concept of roles to:
  - Objects
  - Environments
  - apart from Subjects (entities)
- Primarily targeted towards application writers and administrators
  - Makes policy writing and visualization easier and more intuitive

Michael J. Covington, Matthew J. Moyer and Mustaque Ahamad, "Generalized Role-Based Access Control for Securing Future Applications," *Technical Report GIT-CC-00-02, Georgia Institute of Technology, College of Computing,* February 1, 2000.

Distributed Role-Based Access Control

- No central compliance-checking authority
- Combines RBAC with trust management
- Roles
  - Associated with permissions
  - Augmented by delegations and chains
  - Include delegation rights
  - Valued attributes associated
- Set of delegations stored in a *wallet*
- Compliance with access policy involves
  - Generation of a *proof* (graph of delegations)
  - Distributed credential discovery

E. Freudenthal, T. Pesin, L. Port, E. Keenan and V. Karamcheti, “dRBAC: Distributed Role-Based Access Control for Dynamic Coalition Environments,” *In Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS'02,. IEEE Computer Society*, July 2002.

The “Role” of Trust

- How do we gain and build trust?
  - Identity-based trust
  - Property-based trust
  - Service-based trust
  - Behavior-based trust
    - History of interactions
    - Reputation
- For flexibility, we must separate the process of trust inference from enforcement
- Open questions
  - What are the things I trust ‘X’ with?
  - Does transitive trust really mean anything?

Mechanisms for Building Trust

- Certificates, chains and hierarchies
  - Identity-based trust
- SPKI—Transitive trust through delegation
  - Identity-, Property-, based trust
- Negotiating trust for resource access
  - Potentially subsumes all ways of inferring trust

RFC2693 — C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylonen, “SPKI Certificate Theory.”

Trust Negotiation

- A step towards allowing a system to answer more complex questions:
  - Does ‘X’ possess valid key ‘K’? →
  - Does ‘X’ comply with my access policy ‘P’? →
  - Do we both comply with each others’ policies?
- Primarily targeted towards access control of services on the World Wide Web[1] and the Semantic Web [2]
- Progressive disclosure of trust information (primarily certificate-like credentials), resulting in a yes/no agreement
- Resource and credentials treated alike
  - Governed by individual access control policies
- Is this directly applicable to ubicomp?

[1] M. Winslett, “An Introduction to Trust Negotiation,” *1st International Conference on Trust Management*, Crete, Greece, May 2003.

[2] R. Gavriloaie, W. Nejdl, D. Olmedilla, K. Seamons and M. Winslett, “No Registration Needed: How to Use Declarative Policies and Negotiation to Access Sensitive Resources on the Semantic Web,” *In Proceedings of the 1st First European Semantic Web Symposium*, Heraklion, Greece, May 2004.

Administrative Challenges

- The user can be a liability
- Security-impacting decisions need to be made more frequently
  - The user is not in the best position to make such decisions
  - Much worse than in traditional web-based computing
- How do we provide users with good interfaces
  - To allow them to set and modify policies that will reflect their intentions, and also ensure least violation of security and privacy?
  - That will provide them with understandable feedback about the state of the system?

Discussion Questions

- There are fundamental tradeoffs between security and privacy. How do we reconcile these in pervasive computing?
- How do usability requirements affect security goals?
- What other kinds of security challenges are inherent in pervasive computing environments?
  - Denial of Service (primarily caused by resource constraints)