Security Risks, Low-tech User Interfaces, and Implantable Medical Devices: A Case Study with Insulin Pump Infusion Systems

Nathanael Paul[*]
University of Tennessee
Oak Ridge National Laboratory
pauln@utk.edu

Tadayoshi Kohno[*]
Dept. of Computer Science & Engineering
University of Washington
yoshi@cs.washington.edu

Abstract

Portable implantable medical device systems are playing a larger role in modern health care. Increasing attention is now being given to the wireless control interface of these systems. Our position is that wireless security in portable implantable medical device systems is just a part of the overall system security, and increased attention is needed to address low-tech security issues.

1. Introduction

Recent work in portable implantable medical device systems has highlighted security issues in the wireless control interfaces [Halperin08]. While wireless control security needs to be addressed, other portable implantable medical device (IMD) security challenges remain. We find that user interface (UI) security is under-addressed, and, without requiring technical sophistication, an unauthorized party can significantly harm patients. We consider this area of work an open research problem that needs greater attention.

To understand these issues, we examine insulin pump infusion systems. These systems are complex and require a high amount of user interaction. Any confusion about the current system operation can hurt patient health. User interactions vary based on the pump architecture. In order to interact with the pump device in a "patch pump" architecture, the user uses a remote wireless device that is similar to a smart phone. These control devices store system settings and programs for insulin delivery, and wirelessly transmit commands to the patient's pump device (the transmission range is on the order of meters). The disposable pump device containing the insulin is directly attached to the body, and its subcutaneous insertion point delivers insulin to the body through short plastic tubing. The wireless interface is the main method to control the pump device.

In the more traditional pump architecture, the insulin pump device can be controlled by a physical inter-

Fig. 1: Patch pump remote control display

face on the pump. The pump device has tactile buttons, an electronic display, and it is worn outside of the clothes (i.e., it is not directly attached to the body). Insulin flows from a reservoir inside the pump device through longer tubing (e.g., tens of inches), to an insertion site, and then through a small amount of subcutaneously inserted tubing. While a wireless remote can be added to this architecture (first introduced in 1999), the core architecture involves an interface on the pump itself. Our findings also apply to these systems.

2. Low-tech Security Interface Issues

To decrease the complexity of operating an insulin pump infusion system, the control interface display often hides much of the functionality of the device. The more limited display makes the device easier to use, but the patient trusts that certain settings do not change. In order to change the pump's settings, it is intended that physical possession is needed of the pump remote control or the pump device itself.

Trust in physical possession is unwarranted. Many users set specific settings on the insulin pump, and these settings are changed on a sporadic basis. When some settings are changed that can negatively affect patient health, there is little to nothing in the display to indicate these changes. Opportunities exist to undetectably change device settings since devices are often left unattended during sleep, bathing, or exercise.

Fig. 1 shows a patch pump system that has a wireless control device that is similar to a PDA device. Unlike some wireless interface control security issues, these potential security breaches may have a delayed effect on the patient. For example, current patch pump control devices calculate how much insulin a patient should receive based on patient-specific settings, the current glucose level, and carbohydrate content of a snack or meal. Some patients will simply provide the needed input, and they will not realize if calculations are based on incorrect device settings.

If a device was previously using one unit of insulin per 20 g of carbohydrates as input, a change to one unit of insulin per 10 g of carbohydrates would effectively double the insulin dose during a meal. This setting can be changed in less than 30 seconds on the device shown in Fig. 1, and its effect would be delayed until a bolus correction or bolus change for food. This same issue does exist in a traditional architecture, but its risk

is better mitigated; specifically, in at least one major manufacturer model, the setting can only be changed with the pump device that is tethered to the patient through a long plastic tube that delivers the insulin.

Other similar issues exist in current insulin pump infusion systems. Each pump is designed to deliver a set amount of insulin per unit time (e.g., hour). This amount, called the *basal rate*, can programmatically change according to the device's stored settings. The rate is tuned to each patient in order to maintain euglycemic blood glucose values. If this amount were to undetectably change, this could have the same effect as someone issuing a command to the pump to increase insulin delivery. Of particular note, the amount that is delivered during sleep could be especially important; the patient will need to awake from sleep to address any change that might induce hypoglycemia.

After identifying these issues, we recognize that work is needed both to prevent and detect these events. In prevention, better authentication is needed to stop unwanted changes from occurring. For detection, better user interfaces and improvements in system event recording (i.e., forensics) are needed; audible alerts when a setting is changed are insufficient as frequent alerts can be ignored. Furthermore, more informed user interface may help patients notice a changed setting, but an enhanced UI that alerts a user to a changed critical setting would happen infrequently. The infrequency of the UI alert and its additional complexity would likely result in patient confusion or usage error. The enhanced UI for security has potential safety downsides that suggest a tension between security, safety, and UI design.

In designing better forensics logs, many harmful pump events could be helpful in a different context. Thus, forensics can be improved in the logging of contextual system events. While this will help security, this could also improve patient safety.

3. Towards Detection of Compromise

To improve authentication and system forensics, we are currently investigating how to augment insulin pump components. For example, continuous blood glucose measurement sensors and insertion set devices can potentially protect against low-tech security issues. If these devices can be augmented to add context to system events (e.g., sense when a patient sleeps), the system can make more informed decisions based on the contextual data (e.g., infer that no critical setting should be changed during sleep). For improved authentication, requiring that device control be combined with a patient's sensor or insertion device would mandate that a patient be in close proximity to the control device when it is changed. This authentication requirement should not interfere with an emergency situation where a physician may wish to interact with the insulin pump system.

Conventional approaches for authentication may be useful to enforce proximity, e.g., the use of biometrics [Venk12], but standard trade-offs between safety in emergencies and security also apply. One potential approach to address emergency care needs would be to provide current system data without needing authentication. For instance, certain critical system settings could be displayed at the touch of a physical button. We are currently evaluating and building authentication mechanisms for insulin pump infusion systems.

Similar to these insulin pump systems, both IMDs (e.g., neurostimulators and cardiac devices) and non-implantable, personal medical devices (e.g., CPAP machines) store settings for correct operation. Patients and care givers (e.g., physicians) may assume that previously initialized settings have been left unchanged since initialization, and this trust may cause them to miss a change in a device setting. In some instances, the patient-facing interfaces to some devices may be simpler than that of insulin pump infusion systems, and this may not currently be an issue (e.g., the patient-facing interfaces on a pacemaker are currently minimal, even though the programming machines in a clinic are sophisticated). However, just as insulin pump system interfaces have become more complex, future designs may introduce security risks to the user interfaces.

4. Conclusion

We have detailed security risks in medical device user interfaces. These risks are currently unmitigated in deployed insulin pump infusion systems; similar issues exist in other portable IMDs. Improving the user interface or system event mechanism requires a balancing of device properties, because changes to device design can affect device safety.

Specific augmentation to the continuous blood glucose measurement system, the insertion device, and in the control device may result in improved authentication and in more detailed system events. Through these design changes, low-tech non-wireless security risks can be mitigated.

5. References.

[Halperin08] D. Halperin, et al. Security and Privacy for Implantable Medical Devices. *IEEE Pervasive Computing*. 7(1):30-39. Jan.-Mar. 2008.

[Venk12] K. Venkatasubramanian, A. Banerjee, S. K. S. Gupta. PSKA: Usable and Secure Key Agreement Scheme for Body Area Networks. *IEEE Transactions on Information Technology in Biomedicine*. 14(1):60-68. Jan. 2010.