Towards Context-Aware Cyber-Physical Systems

Radoslav Ivanov, James Weimer, and Insup Lee
PRECISE Center
Computer and Information Science Department
University of Pennsylvania
Email: {rivanov, weimerj, lee}@cis.upenn.edu

I. Introduction

Cyber-Physical Systems (CPS) define a rich class of systems at the intersection of several domains: safety-critical, embedded, real-time, wireless, control, hybrid systems. Multiple CPS crashes and security vulnerabilities over the last few years have exposed the fact that we do not have fail-proof mechanisms for ensuring the safety and security of systems of such complexity. In particular, systems have failed due to different sensor faults (e.g., the recent Tesla crash [2]), actuator faults (e.g., Da Vinci Surgical Systems actuator accidents [1]), and imperfect human-machine interaction (e.g., the crash of Air France Flight 447 off the coast of Brazil [4]). Furthermore, systems have been compromised through sensor attacks [13], [14], communication channel attacks [3] or software vulnerabilities [5].

Due to the complexity of CPS, developing approaches to ensure and verify their safety and security has proven challenging. Standard model-based techniques to verification [12] or anomaly detection [15] cannot be applied because CPS models are rarely known and might often change (e.g., a patient's blood sugar model changes depending on the type of consumed food). At best, limited information is available about the system, e.g., parameterized physiological models (with parameters varying across patients), redundant sensor data, coarse contextual information. In such cases, system designers can develop a combination of techniques utilizing such limited information in order to monitor the system during its operation. In prior work, we developed several such monitoring approaches, namely parameter-invariant (PAIN) anomaly detectors [11], [16] as well as sensor fusion techniques for detecting when the system might be in an unsafe state [9].

In this paper, we provide a brief discussion on using context as a new class of information that can be used to aid estimation/detection approaches to ensuring the safety and security of modern CPS. These systems typically have access to multiple information sources that cannot be directly related to the system state (e.g., mapping oxygen saturation measurements to the overall blood oxygen content is challenging). At the same time, this information is correlated with the system's state (e.g., if the saturation is below a certain threshold, then so is the overall oxygen content). Thus, if formalized properly, context can be used in a way similar to standard measurements.

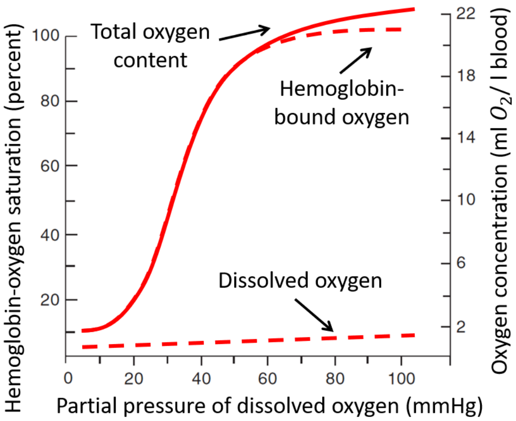

Fig. 1. Most of the oxygen in the blood is bound to hemoglobin. If the hemoglobin-oxygen saturation is below a certain threshold (e.g., 90%), then so is the overall content (e.g., 18 mL oxygen per liter blood).

Although context might provide more coarse information than standard measurements, it can be beneficial from a CPS monitoring point of view, especially in scenarios that are not captured by formal models. Specifically, if the system has access to context that is correlated with unsafe modes of operation, it might be possible to detect such modes and take corrective actions. Thus, context can be used as an extra layer of information in addition to any standard CPS monitoring/testing techniques. The next section provides some intuition as to how to use context, while Sections III and IV present approaches we have developed for using context in estimation [7] and detection [10], respectively.

II. What is Context?

As mentioned in the introduction, context is intuitively defined as information that cannot be directly (functionally) mapped to the system state but that is still correlated with it. An example context measurement can be extracted from hemoglobin-oxygen saturation measurements (available through a pulse oximeter). As illustrated in Figure 1, most of the oxygen in the blood is bound to hemoglobin, so if the hemoglobin-oxygen saturation is below a certain threshold, then the overall blood oxygen content is also below a certain threshold. Thus, the fact that the saturation is below a certain level is a binary context measurement that indicates the state (overall content) is also below a certain level.

A second example of a context measurement is illustrated in Figure 2, which shows data from a typical surgery case at the Children's Hospital of Philadelphia (CHOP). As shown in the figure, missing positive end-expiratory pressure (PEEP) data

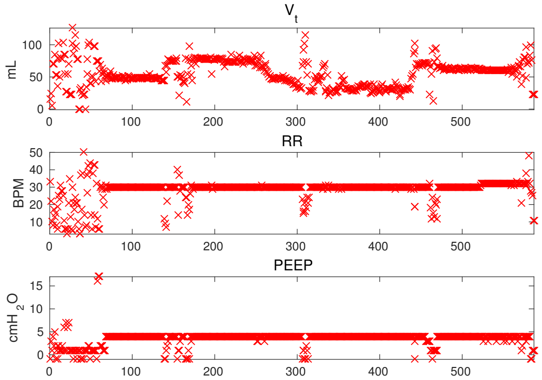

Fig. 2. Typical missing/bad data patterns over time in a surgery case at CHOP. V_t denotes tidal volume (measured in milliliters), RR denotes respiratory rate (measurement in beats per minute, BPM), and PEEP denotes positive end-expiratory pressure (measured in centimeters of water). Missing measurements are set to -1, i.e., they lie on the lower border of each graph.

(an input set by clinicians) are correlated with great noise in the actual measurements. Since all these values are provided by the same machine, one can conclude that the machine is not functioning properly when PEEP is missing (e.g., clinicians have disconnected the machine to check for leaks), i.e., the measurements are meaningless. Thus, missing PEEP can be treated as a context measurement that indicates that available measurements should be treated with caution.

Example context measurements can be found in other domains as well. In an automotive setting, a humidity sensor can be used to detect a fog (i.e., provide a context measurement) such that the system can conclude the camera's measurements might be meaningless. Furthermore, a vehicle can detect nearby buildings using image processing; such a detection (i.e., a context measurement) can be mapped to the vehicle's location since buildings have known positions on the map.

More generally, context can be thought of as discrete data that maps to sets of the state or to system modes (e.g., high- vs. low-variance measurement modes). Context can be given as input (e.g., the patient can indicate their meal type), but it can also be extracted from low-level data using domain expertise (as in Figure 1) or through machine learning techniques (e.g., image processing). Given this intuition, we can use context to aid both estimation and detection problems, as shown next.

III. Context-Aware Estimation

We can see from Figure 1 that some context provides rough information about the system state; thus, it can be used for estimation purposes similar to standard measurements. To capture this intuition, in prior work we modeled context probabilistically given the system state: each binary context measurement has a known probability of being 1 or -1 given the state (e.g., there is a high probability of observing saturation below 90% if the overall oxygen concentration is below 18 mL of oxygen per liter blood). Using this definition, we developed a context-aware filter and provided conditions under which context measurements are sufficient for good estimation [6], [7]. We used the context-aware filter for non-invasive estimation of the patient's blood oxygen content (which cannot be measured non-invasively) and achieved about 20% lower estimation error than prior work [8].

IV. Context-Aware Detection

In addition to estimation, context can also be used to aid detection approaches. In particular, the intuition from Figure 2 can be used to prevent detectors from making wrong decisions during bad-data scenarios. Thus, upon receipt of a context measurement, we treat the actual measurements at that time as having a larger (but unknown) variance. This setting is naturally suited for the PAIN detector mentioned in the introduction [11] as it does not require knowledge of patient-specific parameters, including measurement noise variance. Thus, we developed a context-aware PAIN (CA-PAIN) detector [10] that silences the original PAIN detector during bad-data scenarios, thereby avoiding unnecessary false alarms. The CA-PAIN detector was evaluated both in simulation and on real data, and in both evaluations detected about 5% additional life-critical events without increasing the false alarm rate.

References

[1] Why Intuitive issued a recall for da Vinci surgical system. https://www.advisory.com/daily-briefing/2013/12/06/intuitive-says-da-vinci-surgical-system-can-stall-issues-recall. Accessed: 2016-08-26.

[2] U. N. H. T. S. Administration. Investigation pe 16-007. https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.pdf.

[3] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive experimental analyses of automotive attack surfaces. In *SEC'11: Proc. 20th USENIX conference on Security*, pages 6–6, 2011.

[4] B. d'Enquêtes et d'Analyses. Final report on the accident on 1st June 2009 to the Airbus A330-203 registered F-GZCP operated by Air France flight AF 447 Rio de Janeiro–Paris. *Paris: BEA*, 2012.

[5] N. Falliere, L. O. Murchu, and E. Chien. W32. stuxnet dossier. *White paper, Symantec Corp., Security Response*, 2011.

[6] R. Ivanov, N. Atanasov, M. Pajic, G. Pappas, and I. Lee. Robust estimation using context-aware filtering. In *53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 590–597, 2015.

[7] R. Ivanov, N. Atanasov, M. Pajic, J. Weimer, G. Pappas, and I. Lee. Continuous estimation using context-dependent discrete measurements. In *IEEE Transactions on Automatic Control*. Accepted.

[8] R. Ivanov, N. Atanasov, J. Weimer, M. Pajic, A. Simpao, M. Rehman, G. J. Pappas, and I. Lee. Estimation of blood oxygen content using context-aware filtering. In *Proceedings of the 7th International Conference on Cyber-Physical Systems*, page 28. IEEE Press, 2016.

[9] R. Ivanov, M. Pajic, and I. Lee. Attack-resilient sensor fusion for safety-critical cyber-physical systems. *ACM Trans. Embed. Comput. Syst.*, 15(1):21:1–21:24, Feb. 2016.

[10] R. Ivanov, J. Weimer, and I. Lee. Context-aware detection in medical cyber-physical systems. In *Proceedings of the ACM/IEEE Ninth International Conference on Cyber-Physical Systems (ICCPS)*, 2018. Accepted.

[11] R. Ivanov, J. Weimer, A. F. Simpao, M. A. Rehman, and I. Lee. Prediction of critical pulmonary shunts in infants. *IEEE Transactions on Control Systems Technology*, 24(6):1936–1952, Nov 2016.

[12] M. Pajic, R. Mangharam, O. Sokolsky, D. Arney, J. Goldman, and I. Lee. Model-driven safety analysis of closed-loop medical systems. *IEEE Transactions on Industrial Informatics*, 10(1):3–16, 2014.

[13] S. Peterson and P. Faramarzi. Iran hijacked US drone, says Iranian engineer. *Christian Science Monitor, December*, 15, 2011.

[14] A. H. Rutkin. 'Spoofers' Use Fake GPS Signals to Knock a Yacht Off Course. MIT Technology Review, August 2013.

[15] A. Wald. *Sequential analysis*. Courier Corporation, 1973.

[16] J. Weimer, R. Ivanov, S. Chen, A. Roederer, O. Sokolsky, and I. Lee. Parameter-invariant monitor design for cyber physical systems. *Proceedings of the IEEE*, PP(99):1–22, 2017.