Towards Predictable Wireless Cyber-physical Applications

Octav Chipara and Chenyang Lu
Washington University in St. Louis
{ochipara, lu}@cse.wustl.edu

Abstract

Making wireless cyber-physical applications a reality hinges upon understanding how to build applications with predictable performance. A key challange is to develop communication protocols that enable application components to communicate predictably. We plan on developing a framework that brides wireless communication and real-time schedulability theory under a realistic communication and inference model.

Motivation: Early research on wireless sensor networks focused on applications for which *best-effort* operation is sufficient. Such applications pose limited development challenges since they have low data rates and do not require either reliability or real-time communication. In contrast, the next generation of sensor networks is envisioned to support cyber-physical applications that entail patient monitoring, industrial control, and other critical applications. Today, we are ill-prepared to develop cyber-physical applications due to both theoretical and practical limitations. For example, to develop a patient monitoring application, one must be able to answer whether the sensor network will be able to monitor the vital signs of hundreds of patients reliably. To address this question, current best practice dictates the deployment of the network and the hand tuning of different application parameters (e.g., sampling rates) reaching even in the protocol stack and the operating system (e.g., duty cycle). This process is not only time consuming and ad hoc, but ultimately futile since topology or workload changes may invalidate the tuned parameters.

An application that our work would enable is a sensor network that monitors the vital signs of the patient after their eviction from the intensive care unit. Currently the vital signs of such patients are seldom checked by nurses and manually recorded in their charts. As a result, it may take a long time until a worsening in their condition may be detected, putting their life in danger. Our goal is to reduce the time necessary to detected a worsening in a patient's condition by automating the collection of vital signs and increasing the frequency at which they are collected. This application is challenging because of the following requirements. First, the vital signs of many patients may be routed to the same base-station. Hence, it is important to engineer the application to support *high data rates*. Second, *real-time communication* becomes important when time-critical information (e.g., EKG data) is monitored. Third, the situation of a patient may dynamically change triggering changes in the priority at which the vital signs of the patient should be transmitted. Fourth, the application must provide *reliable data transmission* over multiple hops. Finally, most patients are ambulatory making *mobility* a requirement.

Building such applications is intrinsically difficult due to complex communication and interference relations among nodes. This is compounded by the need to provide spatial reuse for high data rate applications. To remedy this situation, we propose the development of a framework for constructing *predictable* cyber-physical applications by bridging wireless communication and schedulabilitytheory.

Three forms of communication are prevalent in sensor networks: data collection, data dissemination, and data transmission on flows between a node and a base station. Thus, to construct predictable applications it suffices to provide predictable execution for each of the above communication patterns and any of their combinations. We define an operation to be the execution of one of the above communication patterns by the application.

Overview: Our framework has four key components: (1) a communication and interference model, (2) a planner for each communication pattern, (3) a scheduler for executing plans and (4) an admission control component for determining if an application can meet its timing/bandwidth constraints. Next, we describe each component in detail.

Communication and Interference Model: The communication and interference model (ICM) captures our understanding of how a packet transmission affects a set of nodes. Unfortunately, plethora of networking protocols assume overly simplified communication and interference models which results in poor performance and systemic failures in actual deployments. Therefore, an accurate ICM that closely models the reality of radio communication is fundamental for the development of sound protocols.

A good starting point for developing a realistic ICM is the Interference-Communication Graph (ICG) model. The ICG explicitly models the interference and communication relations among nodes as relations in a graph. The graph has all nodes as vertices and has two types of edges: *communication* and *interference* edges. A communication edge indicates that two nodes may communicate reliably while an interference edges capture the interference relations between nodes. The ICG accounts for both asymmetric links and irregular communication ranges which are characteristic of low-power radios. However, the ICG model may be a poor model if either the communication or interference relations vary significantly over time. We plan to empirically validate the ICG model.

Planner: The planner constructs a *plan* for executing a single instance of an operation. A plan is an ordered sequence of *steps*, each comprised of a set of conflict-free transmissions. The planner constructs plans that minimize the time required to perform each operation by carefully accounting for the precedence constraints introduced by each communication pattern.

Scheduler: The scheduler divides time into *slots*. The scheduler runs on every node to determine the slot in which each step in a plan is executed. The scheduler is designed to (1) improve the throughput by executing steps in the plans of multiple operation in the same slot as long as no transmissions in the steps conflict with each other and (2) provide differentiated execution of operations.

Admission Control: The admission control uses our schedulability analysis to determine if an operation meets its deadline. If this is not the case, then the operation is rejected. Note that the admission control component alleviates the need for complex congestion control mechanisms.

The separation between the planner component, which works off-line to construct plans for executing operations, and the scheduler component, which dynamically executes operations based on their temporal properties is the cornerstone of our framework. This setup has intrinsic advantages: (1) We separate the costly operation of constructing plans from their dynamic execution by the scheduler. To reduce the overhead, a key challenge is to reuse a plan to execute multiple operation when it is possible. This is feasible because many operations induce similar communication workloads. For example, a query that computes the average temperature and one that computes maximum humidity in a WSN induce the same workload: each node transmits a single packet to its parent (containing the appropriate information). (2) Unlike most existing TDMA scheduling techniques which construct *static* schedules, in our framework, the scheduler *dynamically* schedules transmissions based on the operation's periods. As a consequence, our approach is significantly more flexible handling workload changes including changes in the periods of operations and their addition/removal: all is required is to disseminate the new/updated rates. However, a key challenge is to construct a computationally efficient local scheduler, which determines what operations to execute at run-time such that no transmission conflicts occur.

Research Issues: To validate our framework we designed a real-time query service that supports data aggregation operations [1]. We also proposed a novel approach for computing upper bounds on response time of data aggregation operation which enables us to perform schedulability analysis and admission control. However, many open challenges still remain including reliability and mobility. Next we sketch possible solutions to these issues.

Achieving reliability in a dynamic system is difficult since many factors may lead to packet losses. Simple solutions such as ARQ or over-provision drastically reduce the throughput of the system. In our framework, the problem is compounded by having to tightly integrate reliability mechanisms with transmission scheduling in order to achieve predictable performance. Since data dissemination and collection are commonly supported by constructed routing trees, we may construct plans that are immune to a class of changes in the routing tree including parent changes and packet drops. To this end, we may modify the routing tree protocol to allow a node to have multiple parents. When a packet is transmitted it is multicast to all its parents. Taking advantage of overhearing, one of the parents that receive the packet correctly will transmit it up the tree.

We are considering two approaches for handling mobility. First, we recognize that patient movement is relatively slow. As a consequence, it may be sufficient to model mobility by addition/removal of edges in the ICG. However, to accomplish this, we must develop more decentralized versions of our algorithms as to facilitate the incremental adjustment of constructed plans as the ICG changes. Second, a common requirement in a hospital is to locate patients. As such an interesting approach would be to combine range-free localization with transmission scheduling. The range-free localization may provide motion profiles for the user, enabling the planner to proactively modify the plans on the fly to handle potential changes in the ICG.

In conclusion, we presented a framework for building predictable cyber-physical applications. The framework features a planner and a scheduler. The planner reduces time to execute an operation by constructing plans that respect the precedence constraints introduced by each communication patterns. The scheduler improves throughput by overlapping the execution of multiple operations and provides differentiated service among operations.

References

[1] O. Chipara, C. Lu, and C.-G. Roman. Real-time query scheduling for wireless sensor networks. In *RTSS'07*.