CPS: Synergy: Resilient Wireless Sensor-Actuator Networks

1. Introduction: Wireless sensor-actuator networks (WSAN) are cyber-physical systems (CPS) consisting of numerous sensing and actuation devices sharing information over a wireless communication network. This project studies resilient WSANs where "resilience" means that the system maintains an awareness of surrounding threats while taking actions assuring a return to operational normalcy as quickly as possible. Building resilient WSANs is challenging due to the time-varying nature of these networks. Temporal variations in a network's quality-of-service introduce an unpredictability that appears to be inconsistent with resilience. Appearances, however, can be deceiving. Wireless networking is evolving rapidly with technologies such as IEEE 802.11n and 3GPP LTE pushing bit/rates from 10^5 bps to 10^8 bps. Moreover, recent results determining the minimum bit-rates for control stability suggest that the actual bandwidth needed for resilience is much less than previously believed. We therefore have two trends, higher achievable bit-rates from wireless technologies and lower required bit-rates for control. Exploiting these two trends will enhance our ability to achieve resilience in large-scale complex WSANs.

We propose exploting these two trends through a three part effort. The first part involves the *development of a hierarchical control architecture* based on a sporadic "event-triggered" message passing framework in which sensors and actuators only exchange information when needed. This approach allows us to greatly reduce the peak and average bit rates required for stabilization and resilience. The second part of this project *develops wireless networking technologies* to support the sporadic bursty traffic generated by our controllers. This work will leverage an emerging paradigm for networking wireless devices known as machine-to-machine (M2M) communication. The third part of this project evaluates the resilience achievable with our approach using a *robotic testbed* consisting of unmanned ground vehicles (UGV) and air vehicles (UAV) that communicate using a USB dongle employing M2M wireless networking technologies.

CPS Relevance: This effort addresses the problem of resilience in CPS. The physical component of our CPS is a multi-robot swarm and the cyber component is the M2M wireless network. The effort addresses the CPS Research Target area *Science of CPS* through its development of fundamental information theoretic limits and practical approaches for resilient control. It addresses the CPS Research Target area *Engineering of CPS* with its development of M2M technologies and layered event-triggered architectures for resilient control. The project's *synergy* comes from its combination of scalable methods in three distinct areas: control, communication, and computer science.

2. Resilient Event-Triggered Control Architecture: A resilient system is one that maintains an active awareness of faults and reacts to faults in a manner that returns the system to normalcy as quickly as possible. The great complexity of large scale WSANs means that anomalies are inevitable, so a major issue in the development of large-scale CPS is their resilience.

Ecological scientists [41] have long recognized that "engineering resilience" as studied in the context of robust stabilization, non-fragile controllers [38], or fault diagnosis and accommodation (FDA) [48] is an inadequate framework for complex networked systems. This is because complex networked systems are susceptible to a wide range of faults arising from human interactions, catastrophic environmental changes, nonlinear feedback interactions leading to regime shifts, failures in sensing or communication, as well as malicious security threats. There is growing awareness of this inadequacy [102] which many feel should be addressed through *hierarchical resilient control systems*.

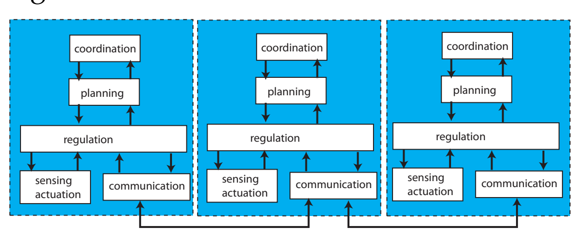

Figure 1: Distributed and Layered WSAN Controller Architecture

Figure 1 shows the proposed hierarchical control architecture. To be concrete, our discussion focuses on an unmanned air vehicle (UAV) swarm, with each shaded box in the figure being a single UAV's controller. This box may be viewed as a *control stack* consisting of a *coordination layer*, a *planning layer*, and a *regulation layer*. The regulation layer is connected directly to the UAV sensors, actuators, and communication interface. This layer regulates the UAV's local state about a known setpoint. The next layer serves a planning function that determines the setpoint for the regulation layer. The setpoint trajectory is obtained by solving an associated receding horizon control (RHC) optimization problem. The RHC cost extremalized by the planning layer is determined by the coordination layer. This layer solves a supervisory control problem that determines the objective that the individual agent is seeking.

Resilience to transient and crash faults are handled in the regulation and coordination layer, respectively. A transient fault may be modeled as an impulsive input causing a discrete jump in a UAV's local state. Section 2.1 presents a novel *event-triggered* method for handling such faults in the regulation layer. Crash faults resulting from the loss of sensor data or communication links are handled in the coordination or planning layer. A supervisory method for assuring resilience to crash faults is described in Section 2.2.

The architecture identified above is similar to that recently proposed in [151]. The differences lie in 1) our use of event-triggered messaging, 2) our separation of the layers concerned with control/coordination and communication, and 3) our focus on coordination and regulation rather than the game-theoretic framework in [151]. We view our effort as being complementary to that in [151] by introducing a more diverse set of tools assuring control resilience.

A novel feature of the proposed work is its use of *event-triggered* feedback streams in which the next measurement is triggered when the plant's output magnitude exceeds a specified threshold. Event-triggering usually generates *sporadic* streams in which the inter-packet time is time-variable and often random, though bounded below by a known constant. Recent experiments have demonstrated that event-triggered control systems use fewer computing and communication resources than time-triggered systems with comparable performance levels. This occurs because a time-triggered system selects its inter-packet interval based on the system's worst-case inputs. Event-triggered systems, on the other hand, adjust the inter-packet time in response to the system's actual state. As a result, event-triggered systems exhibit average inter-packet times that are longer than those seen in comparable time-triggered systems. We believe this feature may help systems exhibit greater resilience to unexpected variations in wireless environment.

Event-triggering may reduce average bit-rates, but it tends to generate sporadic bursty streams of packets. Such bursty flows are not well handled by existing real-time wireless technologies (i.e. Zigbee and wirelessHART). A more appropriate approach would be to use the emerging class of machine-to-machine (M2M) networking technologies [35, 54]. Our research plans for integrating M2M communication with the control architecture in Figure 1 are discussed in Section 3.0.

2.1 - Event-triggered Resilient Regulation: The prevailing wisdom in real-time systems is that periodic time-triggered systems provide the easiest way of assuring the predictability required by safety-critical systems [58, 59]. It may therefore be viewed as foolhardy to use event-triggered streams in safety-critical applications. To justify this approach, consider the system in Figure 2.

Figure 2's system consists of N physical *plants* that are monitored by N sensor/encoder subsystems with each plant being controlled by an actuator/decoder subsystem. The output of the ith plant ($i = 1, 2, \ldots, n$) is a function, $x_i(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$, that satisfies the nonlinear differential equation, $\dot{x}(t) = f_i(x_i(t)) + w_i(t) + u_i(t)$ where f_i is locally Lipschitz and $f_i(0) = 0$. The first input, $w_i(t) = \overline{w}_i \delta(t - T_i)$ is an exogenous impulsive disturbance of unknown magnitude $\overline{w}_i$ that hits the system at unknown time T_i. The other input, $u_i(t)$, is a control signal generated by the ith decoder/actuator subsystem. This signal is an impulse train, $u_i(t) = -\sum_{k=1}^{\infty} \hat{x}_i[k]\delta(t - r_i[k])$ where

$\hat{x}_i[k]$ is the ith encoder's quantized version of the ith plant's output and $r_i[k]$ is the time instant when that quantized measurement was received by the decoder. Define the *quantization map* as the function $Q_i(\cdot) : \mathbb{R} \rightarrow \Omega_i$ where Ω_i is a finite subset of $\mathbb{R}$. Let $\tau_i[k]$ for $k = 1, 2, \ldots, \infty$ denote the kth time instant when the encoder samples the plant's output. The kth consecutive quantized measurement, $\hat{x}_i[k] = Q_i(x_i(\tau_i[k]))$. This measurement is digitally encoded into a packet consisting of $\log_2(|\Omega_i|)$ bits and transmitted to the decoder at time $\tau_i[k]$. The transmission is received at the decoder after a delay $\Delta_i[k] = r_i[k] - \tau_i[k]$.

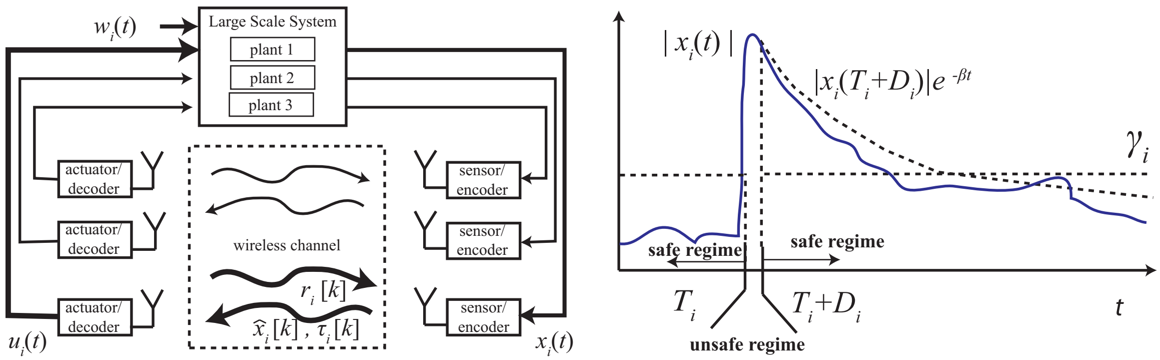

Figure 2: Wireless Sensor Actuator Network (left) - System Performance Specification (right)

The control, $u_i(t)$, acts to reset the plant's output to a neighborhood about zero each time an information packet is received at the decoder. If one can bound the impact of an applied disturbance, w_i, then this control strategy generates actions that are *robust* in the sense that the plant's deviation away from zero is bounded in a predictable manner. We refer to this as the system's *safe operating region* (Figure 2). This strategy, however, presumes the disturbance's impact can be bounded. This will not be be case if the system is hit by a *transient fault* generating an impulse whose magnitude cannot be bounded in an *a priori* manner. Without such a bound, the decoder cannot construct a safe control and the system is uncontrolled over a short period of time, thereby leaving the system in an *unsafe* operating mode as shown in Figure 2.

This project addresses the safety issue by designing controllers that treat safety and robust performance in a unified manner. Systems satisfying these safety and performance requirements are said to be *resilient*. Formally, let T_i denote the fault time. We say the system is *resilient* if there exists a *deadline*, $\overline{D}_i$, such that for any $D_i \leq \overline{D}_i$, one can guarantee

$$|x_i(t)| \leq \max \left\{ |x_i(T_i + D_i)| \, e^{-\beta_i(t - T_i - D_i)}, \gamma_i \right\} \tag{1}$$

for all $t \geq T_i + D_i$ where β_i bounds the *rate* at which the plant's output returns to zero and γ_i is a small positive constant characterizing the desired steady-state deviation from zero. The right side of Figure 2 shows a resilient plant's output where the output increases in an uncontrolled manner after the fault. This uncontrolled behavior continues for duration D_i, after which the decoder has sufficient information to generate a control enforcing the performance bound.

This project's main premise is that system resilience is less costly to maintain using event-triggered streams. The "cost" we refer to is the channel capacity required to support the peak bit-rates generated by the encoder/sensor subsystems. Using methods developed in our earlier papers [136, 133], we can identify the maximum delay that can be tolerated before losing stability. The bit-rate associated with this delay equals the number of bits used to encode $\hat{x}_i[k]$ divided by this maximum delay. Since the plant's output is continuous between consecutive transmissions during the "safe" mode, it can be shown that only a *single bit* is required to encode the information in the sampled state. So the minimum bit-rate assuring the *safe* operation of the plant is simply the reciprocal of the maximum allowable delay. In general, this bit rate is extremely low. For example with a linear system, the rate is proportional to the system's largest unstable eigenvalue.

Let's now examine the bit-rates required during the system's *unsafe* mode. Recall that the system enters its unsafe mode immediately after the impulse $w_i(t)$ hits the system. In this case, the plant's output is no longer a continuous function of time and the plant's output $x_i(t)$ can jump outside of the set defined by the event trigger. This means the decoder's upper bound on the plant's output is invalid and the encoder will need to transmit additional bits to ensure the decoder has enough information to generate a stabilizing control. These extra bits must be transmitted before the safety deadline $\overline{D}_i$ to meet the safety requirement and hence the bit rate generated during the system's unsafe mode will be greater than those rates generated during the system's safe mode.

It should therefore be apparent that the systems in Figure 2 generate two types of bit rates when using event-triggered streams. Event-triggered systems use high bit-rate to resynchronize their encoder and decoder immediately after a fault. They use a much lower bit-rate upon entering their *safe* region. A comparable time-triggered system, on the other hand, uses the same sampling period regardless of whether the system is in its safe or unsafe mode. As a result, time-triggered systems always exhibit greater average channel utilization than event-triggered system. The reason for this difference is simply that event-triggered systems are "situationally aware" of the system's current operating mode and as a result they can switch their channel usage when the system switches between safe and unsafe operational modes.

Does an event-triggered system's lower *average* channel utilization translate into lower infrastructure costs? Communication infrastructure must be sized to handle *peak* loads and these peak loads are the same for both event-triggered and time-triggered streams. For single-user systems, this means there is no benefit to using event-triggered over time-triggered approaches. This will not be true for multi-user systems. For a time-triggered multiple user systems, the channel's capacity will be directly proportional to the number of users. For an event-triggered implementation, however, one only needs to scale channel capacity to the number of simultaneous occurring faults. In general, this will be lower than the total number of users, and so event-triggering reduces infrastructure costs if one *statistically multiplexes* the traffic of multiple users on the channel.

Prior and Preliminary Work: Event-triggered control has been studied in relay [124] and pulse-width modulated [94] systems since the 1960's. Recently this idea has been resurrected as a method to reduce the communication complexity in networked control systems. Threshold based feedback in networked systems and its impact on stability was discussed in [126, 150]. Event-triggering has been suggested as a way to simplify embedded control systems [6] and some analyses have shown that event-triggering reduces computer resource usage [7]. Event triggering approaches based on input-to-state stability [114] and L2 stability [136] have recently appeared.

Prior experimental work [108] established that event-triggering can reduce the average utilization rate in single processor embedded systems without adversely impacting system performance. Experimental work extending these ideas to networked systems [133] demonstrated similar reductions in the average channel usage. This earlier work focused on *increasing the inter-sampling interval* $\tau_i[k+1] - \tau_i[k]$ [4, 5, 139].

This emphasis on inter-sampling interval, in our opinion, is misplaced. Results from [136, 133] show that as one increases the inter-sampling interval, the maximum stabilizing delay, $r_i[k] - \tau_i[k]$, goes to zero. Since the bit-rate generated by an event-triggered system is inversely proportional to this delay, one can see that maximizing the inter-sampling interval results in unbounded bit-rates. Obtaining more realistic estimates of event-triggered bit-rates requires a unification of event-triggering with quantized feedback control [15, 72, 121, 74, 78]. The first steps in this direction have been taken by our group in [69, 68] by establishing conditions under which the stabilizing bit-rates exhibit the efficient attentiveness property [139, 69].

Research Challenges: The preceding discussion made a number of simplifying assumptions that allowed us to focus on the main technical issues regarding event-triggering, bit-rates, and resilience. Our future work will need to remove those restrictions.

- The framework needs to be generalized to loosely coupled systems whose state equations take the form, $\dot{x}_i(t) = f_i(x_i, x_{-i}) + w_i(t) + u_i(t)$ where x_{-i} refers to the local states/outputs of neighboring subsystems in the WSAN. A similar model was studied in [133] and we believe we can use that approach here as well.
- The framework needs to adopt a more realistic disturbance input $w_i(t)$ that includes broadband noise and impulsive disturbances. This can be done by extending our earlier work on broadband disturbances [68].
- The framework shows how one might determine stabilizing bit-rates for event-triggered systems, but it is unclear if these rates are *minimal*. Our analysis in [68] should be "close" to the minimal bit-rate as it relies on many of the same tools used in [15, 121]. So we believe it is possible to establish similar minimal bounds for this problem setup.
- The wireless communication channel's quality-of-service (QoS) is stochastic in nature. The framework should be extended to model channel fading. We propose doing this using a min-plus calculus framework [16] using exponential bounds [145] on the probability of the arriving packets violating the channel's service curve. Recent scaling results [24] suggest this provides a scalable way of evaluating the end-to-end quality of service in multi-hop networks. Our recent work [65] related these bounds to the almost sure stability of simple control systems. We believe this approach provides the correct framework for studying channel fading in event-triggered control systems.
- The communication channel should capture *dropouts* and collisions. Our earlier work [133, 76] found that it was highly desirable to bound the number of consecutively dropped packets. Controlling the number of consecutive dropped messages may be difficult in a real wireless channel. A more realistic position is to encode the quantized measurement as a sequence of packets and then discuss the dropping of packets rather than messages. This encoding should be done so that the loss of a few packets only results in the loss of a few bits, but not the entire quantized measurement. We believe ideas similar to those in [82] can relate these capacity variations back to the system's stochastic stability.

2.2 Resilient Coordination:

Crash faults are sometimes called faults of omission since examples of such faults include the loss of a sensor or the loss of a data link. Since these faults represent a failure of the system, it results in a structural change. If that change makes it impossible for the system to achieve its high level objectives, then we must adaptively re-task the system. This section discusses our proposed approach for resilient cooperative tasking using UAV swarms as an example.

Task Decomposition: It is well known that a collection of robots executing simple reactive actions can display highly complex behaviors in which the entire swarm appears to behave as a single entity. An open question is how one might design these reactive actions to achieve a pre-specified global behavior that is resilient to crash faults.

Our proposed approach is based on our recent work [52] on task decomposition. This prior work proposed a top-down design scheme for multi-robot cooperative tasking through distributed coordination. This approach first assumes that the desired group mission has been specified as a finite state machine (FSM). We then decompose the desired global behaviors into sub-behaviors for each individual robot that when executed by the robot ensure the swarm's global behavior satisfies the specified behavior.

As a simple motivating example, let's consider a two robot coordination problem in which the global task's FSM, A_S, takes the form shown below.

In this figure, the graph's nodes represent the status of the whole robot team. The graph's directed edges represent allowed transitions between theses status. The label, e_i, on the edges represents the ith robot's sensor *event* upon which traversal of the edge is conditioned, say "Robot i enters room". The label, a, represents a coordinated action that is executed by both robots and whose execution again results in a state transition, say "Open the door together". The collection of all events that can be sensed and actuated by the ith robot is denoted as the local event set E_i. In the above example, we have $E_1 = \{a, e_1\}$ and $E_2 = \{a, e_2\}$. It is assumed in [52] that E_i is known *a priori* since one needs to know what each robot is capable of before assigning tasks. It is also assumed that the global desired behavior A_S is defined on the union of E_i. The tasks for Robot i can be obtained by itself in a distributed way through projection A_S (assuming all robots know the global task A_S) onto its own E_i by simply ignoring those events that cannot be observed by the ith robot. The projected subtask for the first robot, denoted as $P_1(A_S)$, is $\longrightarrow \bullet \xrightarrow{e_1} \bullet \xrightarrow{a} \bullet$ and the projection of the second robot is $\longrightarrow \bullet \xrightarrow{e_2} \bullet \xrightarrow{a} \bullet$. If we then have each robot execute the preceding sub-tasks, their group behavior is given by the FSM formed by the parallel composition of $P_1(A_S)$ and $P_2(A_S)$, which is shown in the above figure. This composed automaton is *equivalent* to the original automaton A_S in that they are *bisimilar*. This means that the aggregate behavior obtained when both robots execute their projected (i.e., decomposed) behaviors is equivalent to the originally specified behavior, A_S.

Clearly not all specifications are *decomposable* in this way. For example, if we change the required behavior into A'_S: $\longrightarrow \bullet \xrightarrow{e_1} \bullet \xrightarrow{e_2} \bullet \xrightarrow{a} \bullet$, then it is not decomposable. Actually, A'_S cannot be realized in a distributed way. Results in [52] identify necessary and sufficient conditions for the decomposability of A_S when there are two agents. This work identifies sufficient conditions for decomposability when there are more than two robots. Similar automaton decomposition problems have been studied in the computer science literature, see e.g., [87, 89]. However, to our best knowledge, the automaton decomposition problem still remains open when one considers bisimulation equivalence, as is done here.

Knowing that the original specification, A_S, can be decomposed into the projected $P_i(A_S)$ is important, for these local specifications are much smaller than the original specification. One may, therefore, use recently developed symbolic control methods [115, 117, 9, 22] to construct local reactive control laws for individual robots. Thse synthesized controllers regulate the behavior of individual robots/UAVs at the regulation layer. This project will use these existing synthesis methods for local controller design. The event-triggered methods in section 2.1 provide a resilient implementation of these controllers over a sporadic wireless channel. This section's methods focus on achieving resilient cooperative tasking at the coordination layer (cf. Figure 1).

Resilient Coordination: Robotic swarms are inherently resilient due to the fact that many functions can be performed by more than one robot. If one robot fails, its role can be filled by another robot. Previous simulation and empirical studies have demonstrated this feature nicely, but there are few formal methods to guide the systematic design and deployment of tasks ensuring resilience through this type of redundancy. This project fills that gap by applying the top-down design ideas described above. Our starting point is the assumption that we've already designed controllers for each robot ensuring $P_i(A_S)$ respectively. We consider crash faults arising from the loss of a sensor or a communication link. Such faults cause the event sets E_i to change and so the projected specification,$P_i(A_S)$, also changes. The main question is whether the original task A_S can still be achieved collectively by the team. One way of answering this question in the affirmative is to re-project the specifications and redesign controllers with respect to the new event distributions $\{\tilde{E}_i\}$. This approach, however, is not efficient and it may be difficult to identify the new event set $\{\tilde{E}_i\}$, depending on the nature of the fault detection and identification problem.

We therefore consider a slightly different approach and ask to what extent the originally designed controllers can still achieve the objective without redesign. More precisely, we aim to characterize the fault pattern, i.e. the difference between $\{\tilde{E}_i\}$ and $\{E_i\}$, such that the team behavior after faults is still bisimilar to A_S. The underlying assumption is that these faults can be abstracted as discrete events, say "the color sensor has malfunctioned" or "communication between robot i and j is lost". The consequence of these failures is therefore a change in the elements of E_i. For the loss of a color sensor, one can modify E_i to disregard the color sensor reading. For the loss of a data link, one can revised all events that E_i has in common with E_j into private events. Please note that no global information regarding faults is assumed. Each robot is only aware of the faults around itself and it simply tries to accomplish its previously assigned subtask as best it can. Some robots may not be able to achieve the originally assigned tasks, but as a team, due to the redundancy within the swarm, it may still be possible for the group to collectively enforce A_S. This is, actually, the fault-tolerance issue in top-down design. The exact characterization of tolerable, and in particular intolerable, fault patterns will allow us to pinpoint the fragile points in our multi-robot deployment, and hence help us to introduce redundancies upon demand.

Prior and Related Work: The control of robotic swarms has attracted a great deal of attention. Simulation and empirical studies have shown that a large collection of simple robots executing simple reactive programs can exhibit complex collective behaviors [101, 8, 55, 99]. These studies have generated a great deal of excitement. Recently serious efforts have attempted to develop a rigorous theoretical framework for multi-agent systems. Remarkable efforts have been devoted to the consensus seeking and formation stabilization [90, 50, 119, 99], while approaches like navigation functions [80, 81, 120] and artificial potential functions [60, 103, 97] for distributed formation, optimization-based path planning [111, 98, 88], and game theory-based coordinations [112, 61, 12] have been developed in the literature. A current trend in robotic motion planning is to use formal methods, like model checking and supervisory control, to generate a symbolic path on an abstracted quotient system to satisfy more complicated temporal logic specifications, see e.g., [9, 32, 61, 57, 144, 21]. The critical step, also the most difficult part, of symbolic motion planning is how to obtain an abstraction of the robotic dynamics and environment. Most of research efforts have been devoted to answering the abstraction problem, methods using bisimulation [3, 115, 116] and approximate bisimulation based abstraction [37, 117], maneuver automata [33], and multi-affine control induced workspace partitions [10, 56] have been proposed in the literature. Instead of focusing on a particular abstraction scheme or hybrid controller synthesis scheme, this task leverages off of this prior work and assume that once the specification is given, the lower level motion planning and regulation issues can be efficiently solved by the existing results from the symbolic motion planning literature. This task focuses on obtaining individual specifications whose collective execution achieve the desired global behavior in a resilient manner for swarms of autonomous robots.

Prior work with resilient robotic swarms has focused, primarily, on ground based robots. Much of this work uses empirical studies to assess the fault tolerance of robot swarms using communication [93], market-based strategies [36, 11] and swarm intelligence [143, 23]. Much of this prior work is experimental in nature. The work proposed in this project differs from those earlier studies in that it attempts to provide a formal framework for handling crash faults in robotic swarms based on supervisory control concepts. Another difference is that much of the prior work focused heavily on swarms of unmanned ground vehicles (UGV). This project's long term goal is to test these concepts for UAV swarms very similar to the recent quadrotor swarms at UPenn [85, 86]. To the best of our knowledge, little work has studied the resilient control of quadrotor swarms. Assuring resilience in such swarms is a more difficult problem than what is found in UGV swarms because of the tight coupling that exists between the supervisory and regulation control layers.

The methods proposed in this project are related to the fault-tolerant supervisory control that has been studied in the context of discrete event systems, see e.g., [62, 28, 47, 140]. However, most of the existing results in the fault-tolerant supervisory control literature focus on language specifications and are mainly concerned with the existence of supervisors. This task, on the other hand, focuses on characterizing the faults that a cooperative task is resilient to. It also differs from reliable supervisory control in [118, 79] that seeks the minimal number of supervisors required for correct functionality of the supervised systems. Another related problem is robust supervisory control [73] that designs a supervisor applicable for the whole range of plants.

Proposed Approaches and Research Challenges: The proposed tasks are itemized below.

- In our preliminary study [51], faults are captured as events that can fail to be observed by robots. These faults cause changes in the event distributions $\{\tilde{E}_i\}$. This may be an inadequate strategy for faults that are not directly observable. We therefore propose using techniques in the fault diagnosis literature [49] to identify the faults from observations on event streams.
- Our next concern is whether the global specification A_S is still achievable after the faults are identified and whether the originally designed controllers still maintain safety. Preliminary results obtained in [51] characterize the fault patterns that a specific task A_S can tolerate. In this preliminary study, A_S was assumed to be deterministic while the focus is on the condition that A_S is still decomposable under the new event distributions $\{\tilde{E}_i\}$. We believe that these results can be extended to more general cases and similar results can be obtained to characterize what kinds of fault patterns can be tolerated without redesigning controllers.
- We would like to investigate how the fault-tolerance of multi-agent systems depend on the structural properties of the underlying agent interaction graph. The interaction graph is a connected graph whose nodes are events in E_i and whose edges are labeled by the common events between E_i and its neighbors. This graph reflects the redundancy and information flow between agents. Faults are modeled as the loss of edges in the interaction graph. Graph properties such as connectivity and co-reachability should be related with the multi-agent system's fault-tolerance.
- In the face of faults that cannot be tolerated, the multi-robot system needs to do re-tasking. This redesign should be done in a distributed manner through local reconfigurations. Our preliminary studies have found that turning some private events into public events may make a previously indecomposable A_S become decomposable. We'll investigate methods for addressing this issue by creating new communication links between agents.

3. Machine-to-Machine (M2M) Wireless Networking: Section 2's proposed control architecture employs sporadic message passing and outlines fundamental bounds on the *information rate* rate required for stability and resilience. Older wireless network technologies (.e.g Zigbee and wirelessHART) are poorly suited to sporadic message passing. This task examines the ability of modern wireless networking technologies to support event-triggered control.

The proposed approach to developing resilient WSANs for CPS leverages the emerging class of machine-to-machine (M2M) communications and networking technologies [35, 54]. M2M technologies are building upon the dramatic advances in the commercial wireless industry that have culminated in the recent 3GPP Long Term Evolution (4G LTE) mobile cellular standard and the IEEE 802.11n WiFi wireless local-area network (WiFi) standard. For both, orthogonal frequency division multiplexing (OFDM) with adaptive modulation and channel coding (AMC) and multiple antenna systems (MIMO) is the norm at the physical (PHY) layer. Maximum data rates vary from 10-100 Mbps depending upon the power, bandwidth, range, channel conditions, interference environment, and so forth. Also for both, specific amendments for M2M applications are being developed for applications in healthcare, transportation, smart electric grid, equipment monitoring, intelligent transportation, and so forth. As just one example, there are at least 72 million GSM cellular M2M connections, which is expected to grow to over 280 million by 2016 [100].

It seems natural to forecast that the technology developments and economies of scale anticipated for M2M can fundamentally transform the scope of CPS. Planned M2M amendments to the IEEE 802.11, IEEE 802.16, and 3GPP LTE standards can be viewed as the commercial wireless industry focusing on industrial problems conventionally addressed by ZigBee and WirelessHART technologies. Although the real-time systems and industrial controls communities have made significant progress using these technologies [142, 141, 96, 1], their IEEE 802.15.4 PHY layers provide limited data rate, e.g., maximum 250 kbps per channel, and place stringent limitations on the plant dynamics or number of sensor-actuator nodes in the system. With their significantly higher data rates, adaptive transmission formats, and decreasing costs, the shift to M2M could open up CPS to much larger classes networked control and automation problems.

The opportunity presented by M2M in the broader wireless market is tremendous innovation and investment, but M2M also presents several important research challenges relative to traditional human-to-human (H2H) communications. After carefully selecting a standards family for a given application, closing a real-time control loop around these M2M technologies must address sporadic communication patterns, varying bit rates, latency and resiliency requirements, and potentially a very large number of communicating devices. We believe that the right coupling of M2M technologies with event-triggering and supervisory control as described in Section 2 will provide a viable path for addressing these challenges. In the following sections, we summarize the research directions, preliminary results, and next steps.

3.1: Architecture & Protocol Development: From a communications architecture perspective, strict latency requirements for CPS applications limit packet sizes, which has implications for the rate and reliability of conventional transmission schemes or requires practical implementations of control-specific coding schemes that achieve "anytime reliability" [106, 107]. Sporadic transmissions from event-triggering and moderate packet sizes enable statistical multiplexing among many devices, but at the same time require protocol overhead for synchronization, addressing, multiaccess, and routing to be as small as possible. Many of these issues need to be studied from a fundamental point of view, optimal tradeoffs need to be characterized, efficient practical schemes need to be devised, and insights from these studies need to be merged into emerging standards.

3.1.1: Synchronization for Sporadic Communication. Communication systems are often designed and analyzed assuming continuous transmission of encoded symbols through the channel. However, in many practical applications such as event triggering, this assumption is not valid due to lack of synchronization, shortage of transmission energy, or burstiness of the system. Transmissions become "sporadic" in such scenarios, and the receiver does not explicitly know whether a given channel output results from an encoded message or simply channel noise. Our preliminary results [53] suggest that this form of asynchronous transmission introduces a cost with respect to the achievable rate, which can be interpreted as a form of communications overhead [63].

Another form of asynchronous transmission is modeled in [123] and [132] by a single block transmission that starts at a random time, unknown to the receiver, within an exponentially large window, known to the receiver. In this alternative model, the transmission occurs in one shot and is continuous: once it begins, the whole codeword is transmitted, and before and after transmission the receiver observes only noise. However, in our model for sporadic communication, some number of noise symbols can occur between the symbols of an encoded message, and we consider multiple transmissions inline with the control application.

We now briefly introduce a basic system model, which can easily be extended to apply more generally. A transmitter wants to send a message $M \in \mathcal{M} = \{1, 2, ..., \lceil e^{kR} \rceil\}$ of rate $R = \log |\mathcal{M}|/k$ to a receiver through a discrete memoryless channel (DMC). Let $\mathcal{X}$ and $\mathcal{Y}$ denote the input and output alphabets of the channel, and let $W(y|x)$ denote the channel probability transition matrix. Assume that $\{c^k(m), m = 1, 2, ..., \lceil e^{kR} \rceil\}$ are the codewords utilized by the transmitter, and let X^n and Y^n denote the input and output sequences of the channel, respectively, where $n \geq k$. Here X^n consists of the transmitted codeword for k arbitrary time slots and is equal to a noise symbol denoted by $\star \in \mathcal{X}$ for the other $n - k$ time slots. We define $\alpha := k/n$, where smaller α corresponds to increasingly sporadic communication.

An important first question is what is the capacity [34, 26] of this channel? A simple achievable rate is $C - \frac{1}{\alpha} h(\alpha)$, where $C = \max_P \mathbb{I}(X;Y)$ is the capacity with the maximization over the input distribution $P(x)$, $\mathbb{I}(X;Y)$ is the average mutual information, and $h(\cdot)$ is the binary entropy function [53]. The second term can be interpreted as a penalty on achievable rate caused by the uncertainty in the positions of the encoded symbols at the decoder; however, it turns out that we can improve upon this rate by considering a sophisticated decoding algorithm and analyzing its performance using "partial divergence" [53] that specializes to the well-known Kullback-Leibler divergence as $\alpha \to 1$. The result is that rates less than $\max_P \{\mathbb{I}(X;Y) - f(P, W, \alpha)\}$ are achievable, where $f(\cdot, \cdot, \cdot)$ is an involved functional whose properties are partially explored in [53]. Figure 3 illustrates the two achievable rates (denoted "achv1" and "achv2") for a binary symmetric channel (BSC) with crossover probability p and different values of α.

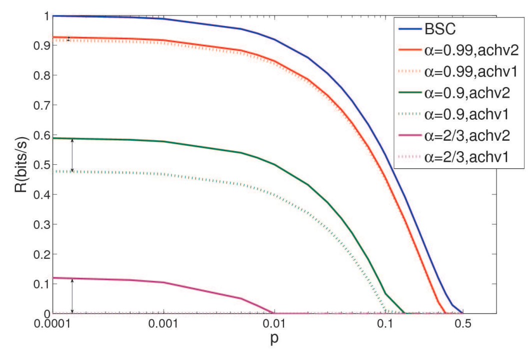

Figure 3: Achievable rate for the sporadic BSC verses cross-over probability p for different α's. As communication becomes more sporadic, i.e., smaller α, synchronization overhead becomes more costly.

Proposed research in this direction will include the following.

- In order to completely characterize the fundamental limits, one needs to develop converse outer bounds for the capacity to complement the achievable inner bounds. Ideally, the inner bound and outer bound meet, and the capacity would be obtained.
- We propose to fully explore the properties of "partial divergence" functional [53], e.g,. convexity, continuity, to provide a deeper understanding of code and protocol design.
- We propose to extend the model and results to two messages at the transmitter, i.e., instead of the noise symbol, the symbols of a codeword containing information about another message is transmitted. Achievability and converse results for the rate region (R_1, R_2), particularly for $R_1 \gg R_2$, would be important for the event-triggering application.
- We plan to design efficient codes and protocols for sporadic communication, e.g., the extent to which separate synchronization "prefixes" are sub-optimal, how self-synchronizing coding schemes can be designed, and so forth. Naturally we will be able to leverage significant

advances in capacity-approaching codes [25] keeping in mind that these codes are typically designed for long, contiguous transmissions.
- Finally, we plan to extend the results to multiaccess scenarios and combine with the models and approaches discussed next.

3.1.2: Multiaccess for Sporadic Communication. Another challenge in M2M networks is to design efficient multiaccess methods to handle a large number of bursty users. Such sporadic multiaccess introduces two practical issues that have not been adequately modeled from a fundamental standpoint: collisions and identification of transmitting users under no collision. We aim to formulate and analyze a model for such scenarios that essentially creates a constrained collision channel model [83] from an arbitrary multiple access channel (MAC) model [26].

For simplicity of exposition, we consider the case of two users. Let $W(y|x_1, x_2)$ be the channel probability law for a discrete memoryless MAC with input alphabets $\mathcal{X}_1$, $\mathcal{X}_2$ and output alphabet $\mathcal{Y}$. If both users transmit simultaneously and we allow sophisticated joint decoding or successive decoding, then the classical MAC model applies, and the transmission rates are constrained by the union of pentagons

$$R_1 < \mathbb{I}(X_1; Y|X_2, T) \quad R_2 < \mathbb{I}(X_2; Y|X_1, T) \quad R_1 + R_2 < \mathbb{I}(X_1, X_2; Y|T) \tag{2}$$

for any input distribution $P(t)P(x_1|t)P(x_2|t)$, where T is a time-sharing random variable [26].

We derive a more practical model from this MAC model by 1) assuming transmissions occur in long slots, 2) allowing the users to operate in either transmitting or sleeping mode in each slot, and 3) forcing the receiver to declare a "collision" rather than jointly decoding if both users transmit, and 4) requiring the receiver to correctly identify and decode the transmitting user under no collision. Sleep is represented by a special input symbol $\star \in \mathcal{X}_1, \mathcal{X}_2$ for each user, and we let $P(x_1)$ on $\mathcal{X}_1$ and $P(x_2)$ on $\mathcal{X}_2$ be two input distributions each with $P(\star) = 0$. The lack of joint decoding eliminates the sum-rate in bound (2), but as we will see introduces other constraints on the individual rates.

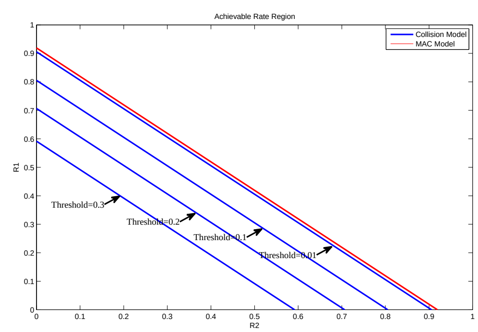

Figure 4: Achievable rates with time-sharing for the sporadic multiaccess collision channel model consisting of BSCs with cross over probability 0.01 under no collision, and a varying lower bound on all state-detection divergence terms. Increasing the divergence threshold to ensure more reliable detection of collisions and identification of the transmitting users reduces the set of achievable transmission rates.

We have analyzed a two-stage receiver that first detects the system "state": no transmission, the first user transmitting, the second user transmitting, or both users transmitting. To each of these states respectively corresponds the following induced output distributions:

$$\begin{aligned} Q_{\star,\star}(y) &= W(y \mid \star, \star) \qquad & Q_{X_1,\star}(y) &= \sum_{x_1 \in \mathcal{X}_1} W(y \mid x_1, \star) P(x_1) \\ Q_{\star,X_2}(y) &= \sum_{x_2 \in \mathcal{X}_2} W(y \mid \star, x_2) P(x_2) \qquad & Q_{X_1,X_2}(y) &= \sum_{\substack{x_1 \in \mathcal{X}_1 \\ x_2 \in \mathcal{X}_2}} W(y \mid x_1, x_2) P(x_1) P(x_2) \end{aligned} \tag{3}$$

A reliable state detector essentially requires that the Kullback-Leibler (KL) divergence [26] between these distributions to be positive. For example, to reliably eliminate Q_{X_1,X_2} (collision) as a possibility if $Q_{X_1,\star}$ (the first user transmitting) prevails, we require $D(Q_{X_1,X_2} \,||\, Q_{X_1,\star}) > t$ for large, but finite, slot lengths, with some threshold $t > 0$. More generally, we would want all pair-wise KL divergences among the distributions (3) to be above some specific threshold in order to ensure the detection error probability decays exponentially in the slot length. In particular, a novel observation is that $D(Q_{X_1,\star} \,||\, Q_{\star,X_2}) > t$, which we expect necessitates that the input distributions of the two users must be distinct in order to identify users in the decoder.

What makes this model interesting is that these divergence constraints translate into constraints on the input distributions $P(x_1)P(x_2)$, which in turn constrains the achievable transmission rates. For example, for the case of only the first user transmitting, reliably received rates must satisfy $R_1 \leq \max_P \mathbb{I}(X_i; Y|\star)$ for $i = 1, 2$ with $P(x_1)$ satisfying the divergence constraints for the state detector, which will be smaller than the individual rate bound in (2). Another view could be that the divergence constraint affects the rate of decay of the decoding error probability, e.g., the error exponent [34]. We note that such tradeoffs are not unlike those in [122, 132].

Proposed research in this direction will include the following.

- We have developed sufficient conditions and reliably achievable rates in our analysis of the two-stage decoder, and we propose to derive corresponding necessary conditions. The latter may require a more complicated, and perhaps more efficient, one-stage decoder.
- We propose to fully develop the tradeoffs among incorrect state detections and decoding errors, using finite slot length tools such as error exponents [34] and dispersion [95, 40] that are extremely important in the context of control applications and CPS.
- We also can apply the same detection scheme as a sensing method for the transmitters to construct an information-theoretic model for carrier sense multiple access (CSMA) and optimize the performance of such schemes.

3.2: Experimental Validation & Standards Enhancements: Given our fundamental understanding of and communications architecture for sporadic communication gained through the activities outlined in Section 3.1, natural but very important questions arise: whether or not resilient WSANs can be constructed using standards-compliant M2M technologies; which underlying M2M technologies are most appropriate for certain CPS applications and to what degree can the communications middleware be application agnostic; and what customizations or extensions of these existing M2M technologies are required in order to make resilient CPS possible based upon event-triggered WSANs.

3.2.1: Assessment and selection among current M2M technologies. Assessment and selection among current M2M-capable technologies will be an important task in building up to the experimental testbed described in Section 4. As we have outlined above, a number of wireless standards "families" are being extended to be more suitable for M2M. We have surveyed these developments, and believe that the IEEE 802.11 family is the best candidate for the proposed project because of its numerous amendments, unlicensed frequencies, and its appropriate range and power levels for the indoor testbed. We expect to acquire or develop a USB dongle built upon an Atheros or Broadcom WiFi chipset and the open8011s software stack (`http://open80211s.org/`), which would allow for specification of the MAC, QoS, and mesh algorithms largely in software.

3.2.2: Experimentally validate communications middleware with selected M2M technologies. We will start with the baseline driver implementations to configure a WSAN for the multi-robot control application and characterize the system stability and performance. We will be able to determine maximum ranges, minimum data rates, and sensitivity to harmful interference for this baseline implementation. With the control application generating its information pattern, i.e., how frequently the event-triggered controller injects packets of certain lengths and rates, we can compute the amount of overhead that the standard IEEE 802.11 protocols introduce and compare to our fundamental performance analyses on overhead for synchronization, multi-access, and routing. We will modify the software MAC layer and adjust PHY layer parameters to emulate our communication schemes and expect to demonstrate some improvements in maximum ranges, minimum data rates, sensitivity to interference, and resiliency.

3.2.3: Propose M2M technology customizations or extensions based upon theoretical and experimental insights. We recognize that there may be constraints imposed by the existing standards and their implementation that may preclude us from demonstrating all of the potential benefits of our communication middleware. If such opportunities arise, we will be in position to make contributions to the evolving M2M standards by demonstrating features that stretch the standards using our flexible software-defined radio platforms, described in the Facilities section.

4. Evaluation Plan for Resilient WSAN Testbed: The resilient control architecture and wireless networking technologies will be evaluated on an indoor multi-robot testbed consisting of a heterogenous mix of UGV's and UAV's communicating over an M2M wireless communication network. The underlying "vision" for the testbed consists of imagining three quadrotors flying as a vertical stack and then simply tossing a ball into the formation to impulsively disturb the formation. Can the methods being developed in the project allow us to recover from such a fault? This level of resilience is, of course, outrageously ambitious, but it is this level of catastrophe that we want to demonstrate resilience against.

There are, of course, numerous multi-robot testbeds at many universities across the country. Examples of UGV testbeds for coordination and control will be found in [104, 27, 84]. A number of fixed wing UAV testbeds have been developed for navigation [31], flight coordination [46], and sensor network monitoring [2]. Rotor UAV testbeds have recently demonstrated swarm coordination [86] and coordination with UGV's [125]. These last two testbeds are most similar to the one being proposed in this project. The novelty in Notre Dame's proposed UAV/UGV testbed rests in its focus on multi-robot swarm resilience and its use of novel M2M communication technologies.

UAV/UGV Testbed Description: The proposed testbed consists of three UGV (ActiveMedia Pioneer robots) and three UAV's (Ascending Technologies Pelican Quadrotor) with the VICON capture system to provide localization information.

The UGV component (see Facilities section) of the testbed was developed under prior NSF funding directed by Dr. Lemmon. It consists of three ActiveMedia Pioneer robots. The Pioneer robot uses acoustic proximity sensors for collision avoidance and gyro-corrected wheel encoders for odometry. It is controlled through an on-board embedded Linux PC that communicates to the Internet over an 802.11 wireless LAN card. Low-level robot motion control is programmed using a set of C++ classes developed by ActivMedia. The vehicles are currently controlled over the Internet using sockets using a remote TCL/TK client. As part of this project, we will modify the wireless networking component of the robots to support the M2M communication technologies. The robot's coordination layer will be modified to realize the supervisory coordination schemes proposed in this project. We will implement event-triggered controllers at the regulation and coordination layers.

For the UAV component of the testbed we propose using the Pelican quadrotor (see Facilities section) from Ascending Technology Company. This platform has a high payload capacity which allows us to mount wireless communication devices designed in this project and a number of sensors. It is designed in a modular fashion allowing the user to change his boards quickly and easily. It is equipped with an Intel Atom processor board (1.6 GHz, 1 GB RAM, 90 g gross weight). We will equip these UAV's with the M2M wireless networking components and rewrite the coordination/regulation controllers using our event-triggered algorithms. The development of the UAV swarm will be directed by Dr. Lin who has been working with similar testbeds at the National University of Singapore since 2006 (see Facilities section).

For indoor mobile robots and UAVs, indoor localization is known to be very important. In algorithm development, it is common practice to assume each robot knows its own position and speed and can obtain its neighbor's position and speed through sensing or communication. Such knowledge, however, is difficult to realize if there is no access to global positioning satellites (GPS). This is, of course, the case for the indoor testbed proposed here. Since the focus of this project is not on indoor localization, we plan to install an "indoor GPS" system based on the well-known commercial VICON motion capture system.

The VICON system has been successfully employed by several UAV groups around the world and has proved ideal for indoor UAV and robot localization. The VICON system monitors light that is reflected off mirrors affixed to the vehicles. The location of these mirrors is captured by the cameras and the system's server assembles this information to calculate speed and position. The VICON system is especially suitable for UAV testbeds because of its fast sampling rate and small time delay. The VICON system will be used to prototype the coordination and control algorithms developed by this project. Initial demos will implement the coordination algorithms on the VICON server. Later demos achieve full autonomy by migrating these algorithms to the individual robots.

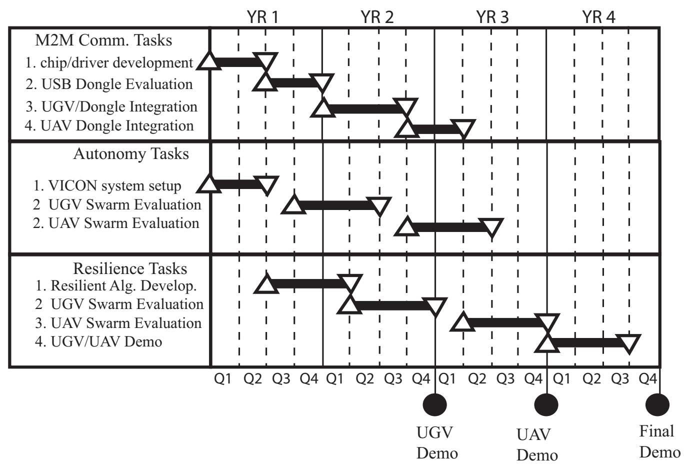

Figure 5: Testbed Development Schedule

The proposed testbed will be used to implement and evaluate the M2M communication technologies described in Section 3 and the resilient control architecture described in Section 2. We've identified three sets of tasks: M2M communication tasks, autonomy tasks, and resilience tasks. These tasks are itemized below. The Gantt chart in Figure 5 shows the proposed schedule with demos planned for the second, third, and fourth year.

- The *M2M Communications task's* objective is to develop a USB dongle for M2M networking that can be plugged into the UGV and UAV platforms. This task, led by Dr. Laneman, has subtasks devoted to selecting communication chipsets and developing the initial chipset drivers. Since integration into the UGV platform will be easier, integration of the dongle with the UGV platform is scheduled for the first half of the second year. Integration of the dongle driver for the UAV platform will be completed by the first quarter of the third year.
- The *autonomy task's* objective is to demonstrate coordinated control of UGV/UAV swarms using time-triggered (periodic) message passing in which all navigation and control functions are executed by the robots. In using periodic message passing we're obtaining a "control" against which the later event-triggered algorithms can be objectively evaluated. This effort will build upon HW/SW developed under Dr. Lemmon and Dr. Lin's earlier projects.
- The *resilience task's* objective is to demonstrate coordinated control of UGV/UAV swarms using event-triggered (sporadic) message passing in which all navigation and control functions are executed by the robots. These tasks will be led by Dr. Lin. Evaluation of the event-triggered UGV platforms should be completed by the end of the second year and evaluation of the event-triggered UAV platforms should be completed by the end of the third year. Assuming these evaluations go well, we intend to do a final combined UGV/UAV demonstration in the fourth year.

5. Intellectual Merit and Broader Impact of Research The intellectual merits of this project rest in its fundamental contributions to supervisory control, event-triggered control, and sporadic communication in M2M networks.

Curriculum Development. The project's testbed will be used to develop a lab-based undergraduate course in CPS. Notre Dame has several graduate CPS courses including "Hybrid dynamical systems", "Formal methods in cyber-physical systems", 'and 'Networked dynamical systems". The PIs will review these offerings in light of the results from this project.

Broader Impacts: The project will broaden its impact through the following activities

- *Industry Collaborations and Technology Transfer.* The PIs have access to a number of programs around Notre Dame that facilitate industrial engagement and technology transfer. The Office of Technology Transfer, Innovation Park startup incubator, and Entrepreneurial MS Program provide, at no cost to the project, staff and students who can help the PIs identify, protect, and commercialize the project's intellectual property. The Wireless Institute has developed a number of corporate relationships, e.g., Sprint, GE Energy, Toyota, and EmNet, and is developing many more relationships with key industry players in the M2M arena. Feedback from these industrial partners will guide us toward addressing real-world problems.
- *Undergraduate Research.* The PIs have supervised many undergraduate students for their summer projects and senior theses. The PIs realize such research opportunities are instrumental in encouraging these students to pursue careers in research and engineering. The experimental aspects of this project provide a natural vehicle for involving undergraduates in research. The PIs will utilize this opportunity through undergraduate summer projects.
- *Outreach and Diversity.* Notre Dame attracts students from minority groups, especially Hispanics. The PIs recognize the importance of exposing high school students to engineering and of broader representation for women and under-represented minority groups. The PIs will engage these students through various campus programs (see Facilities section).

6. Results from Prior NSF Sponsored Research: Dr. Lin has no prior NSF sponsored research. Dr. Lemmon received prior support under NSF grants CNS-0931195 "Dynamically Managing the Real-time Fabric of a Wireless Sensor-Actuator Network" (2009-2012, \$525,000) and ECCS-0925229 "Distributed Optimization, Estimation, and Control of Networked Systems through Event-triggered Message Passing" (2009-2012, \$298,899). Research efforts under these grants investigated the impact of channel burstiness on control system performance [65] and scheduling methods for end-to-end quality of service in wireless networks [42, 45, 44, 146, 43]. Additional work studied event-triggered control [139, 133, 134, 138, 135, 137, 136], estimation [67, 66, 70, 71], and optimization [131, 129, 130, 127, 128]. These grants studied the relation between quantized and event-triggered control [68, 69, 76, 75, 78, 77]. A tutorial overview will be found in [64]. We note that after no-cost extensions these projects are scheduled to be completed on August 31, 2013.

Dr. Laneman has received prior support under NSF grant CNS-06-26595, *Collaborative Research: NeTS-ProWin-NBD: A New Taxonomy for Cooperative Wireless Networking* (2006–2012, \$441,824). This project studied relaying and cooperative communications which led to new link abstractions for wireless networks. The PI's efforts for this collaborative project followed the main areas: high-level architectural constructs that involve relaying, specific relay processing algorithms and protocols, development of a network testbed of software-defined radios (SDRs) for prototyping and refinement of algorithms and architectures. Publications resulting from this project include [19, 20, 18, 109, 110, 92, 91, 17, 148, 149, 147, 113, 105, 13, 14, 39, 30, 29], and we note that after no-cost extensions the project is schedule to be completed on August 31, 2012.

References

[1] J. Akerberg, M. Gidlund, and M. Bjorkman. Future Research Challenges in Wireless Sensor and Actuator Networks Trageting Industrial Automation. In *Proc. IEEE Int. Conf. Industrial Informatics*, pages 410–415, July 2011.

[2] J. Allred, A.B. Hasan, S. Panichsakul, W. Pisano, P. Gray, J. Huang, R. Han, D. Lawrence, and K. Mohseni. Sensorflock: an airborne wireless sensor network of micro-air vehicles. In *Proceedings of the 5th international conference on Embedded networked sensor systems*, pages 117–129. ACM, 2007.

[3] R. Alur, T.A. Henzinger, G. Lafferriere, and G.J. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(7):971 –984, July 2000.

[4] A. Anta and P. Tabuada. Isochronous manifolds in self-triggered control. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 3194–3199. IEEE, 2009.

[5] A. Anta and P. Tabuada. On the minimum attention and anytime attention problems for nonlinear systems. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 3234–3239. IEEE, 2010.

[6] K.E. Arzen. A simple event-based pid controller. In *Proceedings of the 14th IFAC World Congress*, 1999.

[7] K.J. Astrom and B.M. Bernhardsson. Comparison of riemann and lebesgue sampling for first order stochastic systems. In *Proceedings of the IEEE Conference on Decision and Control*, 1999.

[8] T. Balch and R.C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1994.

[9] C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, and G.J. Pappas. Symbolic planning and control of robot motion. *Robotics and Automation Magazine, IEEE*, 14(1):61–70, March 2007.

[10] C. Belta, V. Isler, and G.J. Pappas. Discrete abstractions for robot motion planning and control in polygonal environments. *Robotics, IEEE Transactions on*, 21(5):864 – 874, oct. 2005.

[11] M. Bernardine Dias, M. Zinck, R. Zlot, and A. Stentz. Robust multirobot coordination in dynamic environments. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 4, pages 3435–3442. IEEE, 2004.

[12] Sourabh Bhattacharya, Tamer Basar, and Naira Hovakimyan. Singular surfaces in multi-agent connectivity maintenance games. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 261 –266, dec. 2011.

[13] Glenn J. Bradford and J. Nicholas Laneman. A Survey of Implementation Efforts and Experimental Design for Cooperative Communications. In *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Dallas, TX, March 2010. Invited paper.

[14] Glenn J. Bradford and J. Nicholas Laneman. Low Latency Relaying Schemes for Next-Generation Cellular Networks. In *Proc. IEEE Int. Commun. Conf. (ICC)*, Ottawa, Canada, June 2012. To appear.

[15] R.W. Brockett and D. Liberzon. Quantized feedback stabilization of linear systems. *Automatic Control, IEEE Transactions on*, 45(7):1279–1289, 2000.

[16] A. Burchard, J. Liebeherr, and S.D. Patek. A min-plus calculus for end-to-end statistical service guarantees. *Information Theory, IEEE Transactions on*, 52(9):4105–4114, 2006.

[17] Woohyuk Chang, Shivaprasad Kotagiri, J. Nicholas Laneman, Sae-Young Chung, and Yong-Hoon Lee. Compress-Forward Relaying over Parallel Gaussian Channels. In *Proc. Comp. Adv. Mult-Sensor and Adaptive Processing*, St. Thomas, US Virgin Islands, December 2007.

[18] Deqiang Chen, Kambiz Azarian, and J. Nicholas Laneman. A Case for Amplify-Forward Relaying in the Block-Fading Multiaccess Channel. *IEEE Trans. Inform. Theory*, 54(8):3728–3733, August 2008.

[19] Deqiang Chen, Martin Haenggi, and J. Nicholas Laneman. Distributed Spectrum-Efficient Routing Algorithms in Wireless Networks. In *Proc. Conf. Inform. Sci. and Syst. (CISS)*, Baltimore, MD, March 2007.

[20] Deqiang Chen, Martin Haenggi, and J. Nicholas Laneman. Distributed Spectrum-Efficient Routing Algorithms in Wireless Networks. *IEEE Trans. Wireless Commun.*, 7(12):5297 – 5305, December 2008.

[21] Y. Chen, X. Ding, A. Stefanescu, and C. Belta. Formal approach to the deployment of distributed robotic teams. *Robotics, IEEE Transactions on*, 28(1):158 –171, feb. 2012.

[22] Y. Chen, X. C. Ding, A. Stefanescu, and C. Belta. A formal approach to deployment of robotic teams in an urban-like environment. In *10th International Symposium on Distributed Autonomous Robotics Systems (DARS), Lausanne, Switzerland, 2010*, 2010.

[23] A.L. Christensen, R. OGrady, and M. Dorigo. From fireflies to fault-tolerant swarms of robots. *Evolutionary Computation, IEEE Transactions on*, 13(4):754–766, 2009.

[24] F. Ciucu, A. Burchard, and J. Liebeherr. Scaling properties of statistical end-to-end bounds in the network calculus. *Information Theory, IEEE Transactions on*, 52(6):2300–2312, 2006.

[25] Daniel J. Costello, Jr. and G. David Forney, Jr. Channel Coding: The Road to Channel Capacity. *Proc. IEEE*, November 2006. Submitted for publication.

[26] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., New York, 1991.

[27] D. Cruz, J. McCLINTOCK, B. Perteet, O.A.A. Orqueda, Y. Cao, and R. Fierro. Decentralized cooperative control-a multivehicle platform for research in networked embedded systems. *Control Systems, IEEE*, 27(3):58–78, 2007.

[28] H. Darabi, M.A. Jafari, and A.L. Buczak. A control switching theory for supervisory control of discrete event systems. *Robotics and Automation, IEEE Transactions on*, 19(1):131 – 137, February 2003.

[29] Michael L. Dickens, Brian P. Dunn, and J. Nicholas Laneman. Design and Implementation of a Portable Software Radio. *IEEE Commun. Mag.*, 46(8):58–66, August 2008.

[30] Michael L. Dickens, Brian P. Dunn, and J. Nicholas Laneman. Portable Software Radios using Commodity Hardware and Open-Source Software. In *Proc. SDR*, Washington, DC, October 2008.

[31] J. Evans, G. Inalhan, J.S. Jang, R. Teo, and C.J. Tomlin. Dragonfly: A versatile uav platform for the advancement of aircraft navigation and control. In *Digital Avionics Systems, 2001. DASC. The 20th Conference*, volume 1, pages 1C3–1. IEEE, 2001.

[32] Georgios E. Fainekos, Antoine Girard, Hadas Kress-Gazit, and George J. Pappas. Temporal logic motion planning for dynamic robots. *Automatica*, 45(2):343–352, 2009.

[33] E. Frazzoli, M. A. Dahleh, and E. Feron. Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE Trans. on Robotics*, 21(6):1077C1091, 2005.

[34] Robert G. Gallager. *Information Theory and Reliable Communication*. John Wiley & Sons, Inc., New York, 1968.

[35] Wu Geng, S. Talwar, K. Johnsson, N. Himayat, and K.D. Johnson. M2M: From Mobile to Embedded Internet. *IEEE Comm. Mag.*, 49(4):36–43, April 2011.

[36] B.P. Gerkey and M.J. Mataric. A market-based formulation of sensor-actuator network coordination. In *Proceedings of the AAAI Spring Symposium on Intelligent Embedded and Distributed Systems*, pages 21–26, 2002.

[37] Antoine Girard and George J. Pappas. Hierarchical control system design using approximate simulation. *Automatica*, 45(2):566–571, 2009.

[38] W.M. Haddad and J.R. Corrado. Resilient dynamic controller design via quadratic lyapunov bounds. In *Decision and Control, 1997., Proceedings of the 36th IEEE Conference on*, volume 3, pages 2678–2683. IEEE, 1997.

[39] Eric Hardouin, J. Nicholas Laneman, Suzuki Hidetoshi, Alex Golitschek, and Osvaldo Gonsa. *LTE - The UMTS Long Term Evolution: From Theory to Practice*, chapter Relaying. John Wiley and Sons, Second edition, 2011.

[40] M. Hayashi. Information Spectrum Approach to Second-Order Coding Rate in Channel Coding. *IEEE Trans. Information Theory*, 55(11):4947–4966, nov 2009.

[41] C.S. Holling. Engineering resilience versus ecological resilience. In P.C. Schulze, editor, *Engineering within ecological constraints*, pages 31–44. National Academy Press, Washington D.C., 1996.

[42] S. Hong, T. Chantem, and X.S. Hu. Meeting end-to-end deadlines through distrubted deadline assignment. In *IEEE Real-time systems Symposium (RTSS)*, Vienna, Austria, 2011.

[43] S. Hong, X. Hu, and M.D. Lemmon. An adaptive approach to reduce control delay variations. In *Real-time Systems Symposium - work-in-progress session*, 2009.

[44] S. Hong, X.S. Hu, and M.D. Lemmon. Reducing delay jitter of real-time control tasks through adaptive deadline adjustments. In *Euromicro conference on real-time systems (ECRTS10)*, 2010.

[45] S. Hong, X.S. Hu, and M.D. Lemmon. An adaptive transmission rate control approach to minimize energy consumption. In *IEEE Real-time and Embedded Technology and Applications Symposium (RTAS)*, Chicago, IL, USA, 2011.

[46] J. How, E. King, and Y. Kuwata. Flight demonstrations of cooperative control for uav teams. In *AIAA 3rd" Unmanned Unlimited" Technical Conference, Workshop and Exhibit*, pages 20–23, 2004.

[47] M. V. Iordache and P. J. Antsaklis. Resilience to failure and reconfigurations in the supervision based on place invariants. In *Proceedings of the 2004 American Control Conference*, pages 4477–4482, 2004.

[48] R. Isermann. Process fault detection based on modeling and estimation methods–a survey. *Automatica*, 20(4):387–404, 1984.

[49] R. Isermann. Supervision, fault-detection and fault-diagnosis methods–an introduction. *Control engineering practice*, 5(5):639C652, 1997.

[50] A. Jadbabaie, Jie Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *Automatic Control, IEEE Transactions on*, 48(6):988–1001, June 2003.

[51] M. Karimadini and H. Lin. Fault-tolerant cooperative tasking for multi-agent systems. *Internationl Journal of Control*, 84(12):2092–2107, 2011.

[52] M. Karimadini and H. Lin. Guaranteed global performance through local coordinations. *Automatica*, 47(5):890–898, 2011.

[53] Mostafa Khoshnevisan and J. Nicholas Laneman. Achievable Rates for Intermittent Communication. In *Proc. IEEE Int. Symp. Information Theory (ISIT)*, Boston, MA, July 2012. Submitted for publication.

[54] Chang Kim, A. Soong, M. Tseng, and Xiang Zhixian. Global Wireless Machine-to-Machine Standardization. *IEEE Internet Computing*, 15(2):64–69, April 2011.

[55] M. Kloetzer and C. Belta. Temporal logic planning and control of robotic swarms by hierarchical abstractions. *Robotics, IEEE Transactions on*, 23(2):320–330, April 2007.

[56] M. Kloetzer and C. Belta. A fully automated framework for control of linear systems from temporal logic specifications. *Automatic Control, IEEE Transactions on*, 53(1):287–297, Feb. 2008.

[57] M. Kloetzer and C. Belta. Automatic deployment of distributed teams of robots from temporal logic motion specifications. *Robotics, IEEE Transactions on*, 26(1):48 –61, feb. 2010.

[58] H. Kopetz. *Real-time systems: design principles for distributed embedded applications*, volume 25. Springer-Verlag New York Inc, 2011.

[59] H. Kopetz and G. Grunsteidl. Ttp-a time-triggered protocol for fault-tolerant real-time systems. In *Fault-Tolerant Computing, 1993. FTCS-23. Digest of Papers., The Twenty-Third International Symposium on*, pages 524–533. IEEE, 1993.

[60] Y. Koren and J. Borenstein. Potential field methods and their inherent limitations for mobile robot navigation. In *1991 IEEE International Conference on Robotics and Automation*, pages 1398 – 1404. IEEE, 1991.

[61] H. Kress-Gazit, G.E. Fainekos, and G.J. Pappas. Temporal-logic-based reactive mission and motion planning. *Robotics, IEEE Transactions on*, 25(6):1370 –1381, 2009.

[62] S. Lafortune and F. Lin. On tolerable and desirable behaviors in supervisory control of discrete event systems. In *Decision and Control, 1990., Proceedings of the 29th IEEE Conference on*, pages 3434 –3439 vol.6, December 1990.

[63] J. Nicholas Laneman and Brian P. Dunn. Communications Overhead as the Cost of Constraints. In *Proc. IEEE Information Theory Workshop (ITW)*, Paraty, Brazil, October 2011.

[64] M.D. Lemmon. Event-triggered feedback in control, estimation, and optimization. In A. Bemporad, M. Heemels, and M. Johansson, editors, *Networked Control Systems*, volume 405 of *Lecture Notes in Control and Information Sciences*, pages 293–358. Springer Verlag, 2010.

[65] M.D. Lemmon and X.S. Hu. Almost sure stability of networked control systems under exponentially bounded bursts of dropouts. In *Hybrid Systems: computation and control*, Chicago, IL, USA, 2011.

[66] L. Li and M. Lemmon. Performance and average sampling period of sub-optimal triggering event in event triggered state estimation. In *IEEE Conference on Decision and Control*, 2011.

[67] L. Li and M. Lemmon. Weakly coupled event triggered output feedback control in wireless networked control systems. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*, pages 572–579. IEEE, 2011.

[68] L. Li, X. Wang, and M.D. Lemmon. Stabilizing bit-rate of perturbed event-triggered control systems. In *Proceedings of the 4th IFAC Conference on Analysis and Design of Hybrid Systems*, Eindhoven, The Netherlands, 2012.

[69] L. Li, X. Wang, and M.D. Lemmon. Stabilizing bit-rates in quantized event triggered control systems. In *Hybrid Systems: computation and control*. Beijing, China, 2012.

[70] Lichun Li and M.D. Lemmon. Event-triggered output feedback control of finite horizon discrete-time multi-dimensional linear processes. In *Proceedings of the IEEE Conference on Decision and Control*, 2010.

[71] Lichun Li, M.D. Lemmon, and X. Wang. Event-triggered state estimation in vector linear processes. In *Proceedings of the American Control Conference*, 2010.

[72] D. Liberzon. Nonlinear stabilization by hybrid quantized feedback. *Hybrid Systems: Computation and Control*, pages 243–257, 2000.

[73] F. Lin. Robust and adaptive supervisory control of discrete event systems. *Automatic Control, IEEE Transactions on*, 38(12):1848 –1852, December 1993.

[74] Q. Ling and M.D. Lemmon. Stability of quantized control systems under dynamic bit assignment. *IEEE Transactions on Automatic Control*, 50(5):734–740, 2005.

[75] Q. Ling and M.D. Lemmon. Input-to-state stabilizability of quantized linear control systems under feedback dropouts. In *Proceedings of the American Control Conference*, 2010.

[76] Q. Ling and M.D. Lemmon. A necessary and sufficient feedback dropout condition to stabilize quantized linear control systems with bounded noise. *IEEE Transactions on Automatic Control*, 55(6):1494–1500, 2010.

[77] Q. Ling, M.D. Lemmon, and H. Lin. Stabilize and n-dimensional quantized nonlinear feedforward system with 1 bit. In *IEEE Conference on Decision and Control*, 2009.

[78] Q. Ling, M.D. Lemmon, and H. Lin. Asymptotic stabilization of dynamically quantized nonlinear systems in feedforward form. *Journal of Control Theory and Applications*, 8(1):27–33, 2010.

[79] Fuchun Liu and Hai Lin. Reliable supervisory control for general architecture of decentralized discrete event systems. *Automatica*, 46(9):1510 – 1516, 2010.

[80] S. G. Loizou and K. J. Kyriakopoulos. *Multirobot Navigation Functions I, Book Series Studies in in H. A. P Blom and J. Lygeros (Eds.), Stochastic Hybrid Systems: Theory and Safety Critical Applications*. Springer, 2006.

[81] S. G. Loizou and K. J. Kyriakopoulos. *Multirobot Navigation Functions II, Book Series Studies in in H. A. P Blom and J. Lygeros (Eds.), Stochastic Hybrid Systems: Theory and Safety Critical Applications*. Springer, 2006.

[82] N.C. Martins, M.A. Dahleh, and N. Elia. Feedback stabilization of uncertain systems in the presence of a direct link. *Automatic Control, IEEE Transactions on*, 51(3):438–447, 2006.

[83] J. Massey and P. Mathys. The Collision Channel without Feedback. *IEEE Trans. Information Theory*, 31(2):192–204, mar 1985.

[84] N. Michael, J. Fink, and V. Kumar. Experimental testbed for large multirobot teams. *Robotics & Automation Magazine, IEEE*, 15(1):53–61, 2008.

[85] N. Michael, J. Fink, and V. Kumar. Cooperative manipulation and transportation with aerial robots. *Autonomous Robots*, pages 1–14, 2009.

[86] N. Michael, D. Mellinger, Q. Lindsey, and V. Kumar. The grasp multiple micro-uav testbed. *Robotics & Automation Magazine, IEEE*, 17(3):56–65, 2010.

[87] Rémi Morin. Decompositions of asynchronous systems. In *CONCUR '98: Proceedings of the 9th International Conference on Concurrency Theory*, pages 549–564, London, UK, 1998. Springer-Verlag.

[88] N. Motee, A. Jadbabaie, and G. Pappas. A duality approach to path planning for multiple robots. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 935 – 940, May 2010.

[89] M. Mukund. *From global specifications to distributed implementations, in B. Caillaud, P. Darondeau, L. Lavagno (Eds.),* Synthesis and Control of Discrete Event Systems*, Kluwer*. Springer Berlin / Heidelberg, pp. 19–35, Berlin, 2002.

[90] R. Olfati-Saber and R.M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *Automatic Control, IEEE Transactions on*, 49(9):1520 – 1533, sept. 2004.

[91] Özgür Oyman and J. Nicholas Laneman. Multihop Diversity in Wideband OFDM Systems: The Impact of Spatial Resuse and Frequency Selectivity. In *Proc. Int. Symp. Spread Spectrum Techniques and Applications (ISSSTA)*, Bologna, Italy, August 2008.

[92] Özgür Oyman, J. Nicholas Laneman, and Sumeet Sandhu. Multihop Relaying for Broadband Wireless Mesh Networks: From Theory to Practice. *IEEE Commun. Mag.*, 45(11):116–122, November 2007.

[93] L.E. Parker. Alliance: An architecture for fault tolerant multirobot cooperation. *Robotics and Automation, IEEE Transactions on*, 14(2):220–240, 1998.

[94] E. Polak. Stability and graphical analysis of first order of pulse-width modulated sampled data regulator systems. *IRE Trans. Automatic Control*, AC-6(3):276–282, 1963.

[95] Y. Polyanskiy, H. V. Poor, and Sergio Verdú. Channel Coding Rate in the Finite Blocklength Regime. *IEEE Trans. Information Theory*, 56(5):2307–2359, may 2010.

[96] J. Proenza and S. Vitturi. Guest Editorial Special Section on Industrial Communication Systems. *IEEE Trans. Industrial Informatics*, 6(3):365–368, August 2010.

[97] J. H. Reif and H. Wang. *Social potential fields: A distributed behavioral control for autonomous robots, In K. Goldberg, D. Halperin, J. C. Latombe, R. Wilson, (Eds.),* The Algorithmic Foundations of Robotics, volume 8. A. K. Peters, Boston, MA, 1995.

[98] C. Reinl and O. v. Stryk. Optimal control of multi-vehicle-systems under communication constraints using mixed-integer linear programming. In *RoboComm '07: Proceedings of the 1st international conference on Robot communication and coordination*, pages 1–8, Piscataway, NJ, USA, 2007. IEEE Press.

[99] W. Ren and R. W. Beard. *Distributed Consensus in Multi-Vehicle Cooperative Control*. Springer Verlag, New York, 2007.

[100] Pyramid Research. The Machine-to-Machine Market: A High-Growth Opportunity for MNOs, October 2011.

[101] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34, New York, NY, USA, 1987. ACM.

[102] C.G. Rieger, D.I. Gertman, and M.A. McQueen. Resilient control systems: next generation design research. In *Human System Interactions, 2009. HSI'09. 2nd Conference on*, pages 632–636. IEEE, 2009.

[103] E. Rimon and D.E. Koditschek. Exact robot navigation using artificial potential functions. *Robotics and Automation, IEEE Transactions on*, 8(5):501 – 518, 1992.

[104] CL Robinson, HJ Schutz, G. Baliga, and PR Kumar. Architecture and algorithm for a laboratory vehicle collision avoidance system. In *Intelligent Control, 2007. ISIC 2007. IEEE 22nd International Symposium on*, pages 23–28. IEEE, 2007.

[105] Peter Rost, Gerhard Fettweis, and J. Nicholas Laneman. Opportunities, Constraints, and Benefits of Relaying in the Presence of Interference. In *Proc. IEEE Int. Commun. Conf. (ICC)*, Dresden, Germany, June 2009.

[106] Anant Sahai and Sanjor Mitter. The Necessity and Sufficiency of Anytime Capacity for Stabilization of a Linear System over a Noisy Communication Link, Part I: Scalar Systems. *IEEE Trans. Information Theory*, 52(8):3369–3395, aug 2006.

[107] Anant Sahai and Sanjor Mitter. The Necessity and Sufficiency of Anytime Capacity for Stabilization of a Linear System over a Noisy Communication link, Part II: Vector Systems. Submitted for publication., dec 2007.

[108] J. Sandee, W. Heemels, and P. van den Bosch. Case studies in event-driven control. *Hybrid Systems: computation and control*, pages 762–765, 2007.

[109] Anna Scaglione, Dennis Goeckel, and J. Nicholas Laneman. Cooperative Communications in Mobile Ad-Hoc Networks: Rethinking the Link Abstraction. *IEEE Signal Processing Mag.*, 23(5):18–29, September 2006.

[110] Anna Scaglione, Dennis Goeckel, and J. Nicholas Laneman. *Distributed Antenna Systems: Open Architecture for Future Wireless Communications*, chapter Cooperative Communications in Mobile Ad-Hoc Networks: Rethinking the Link Abstraction, pages 87–116. Auerbach Publications, CRC Press, 2007.

[111] T. Schouwenaars, B. DeMoor, E. Feron, and J. How. Mixed integer programming for multi-vehicle path planning. In *European control conference, ECC 2001*, pages 2603–2608, 2001.

[112] E. Semsar-Kazerooni and K. Khorasani. Multi-agent team cooperation: A game theory approach. *Automatica*, 45(10):2205–2213, 2009.

[113] Harold Sneessens, Luc Vandendorpe, and J. Nicholas Laneman. Adaptive CF Relaying in Fading Environments with or without Wyner-Ziv Coding. In *Proc. IEEE Int. Commun. Conf. (ICC)*, Dresden, Germany, June 2009.

[114] P. Tabuada. Event-triggered real-time scheduling of stabilizing control tasks. *Automatic Control, IEEE Transactions on*, 52(9):1680–1685, 2007.

[115] P. Tabuada and G.J. Pappas. From discrete specifications to hybrid control. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 4, pages 3366–3371 vol.4, Dec. 2003.

[116] P. Tabuada and G.J. Pappas. Linear time logic control of discrete-time linear systems. *Automatic Control, IEEE Transactions on*, 51(12):1862–1877, Dec. 2006.

[117] Paulo Tabuada. *Verification and Control of Hybrid Systems: A Symbolic Approach*. Springer, Springer New York, 2009.

[118] S. Takai and T. Ushio. Reliable decentralized supervisory control of discrete event systems. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 30(5):661 –667, October 2000.

[119] H.G. Tanner, A. Jadbabaie, and G.J. Pappas. Stable flocking of mobile agents part ii: dynamic topology. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 2, pages 2016 – 2021 Vol.2, dec. 2003.

[120] H.G. Tanner and A. Kumar. Towards decentralization of multi-robot navigation functions. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 4132 – 4137, april 2005.

[121] S. Tatikonda and S.K. Mitter. Control under communication constraints. *IEEE Transactions on Automatic Control*, 49(7):1056–1068, 2004.

[122] A. Tchamkerten, V. Chandar, and G. Wornell. Communication under Strong Asynchronism. *IEEE Trans. Information Theory*, 55(55):4508–4528, October 2009.

[123] A. Tchamkerten, V. Chandar, and G. Wornell. Asynchronous Communication: Capacity Bounds and Suboptimality of Training. Submitted for publication., May 2011.

[124] Y.Z. Tsypkin. *Relay Control Systems*. Cambridge University Press, 1984.

[125] P. Twu, R. Chipalkatty, J.P. de la Croix, T. Ramachandran, J. Shively, M. Egerstedt, A. Rahmani, and R. Young. A hardware testbed for multi-uav collaborative ground convoy protection in dynamic environments. In *Proceedings AIAA Modeling and Simulation Technologies Conference*, Portland, OR, 2011.

[126] G.C. Walsh, H. Ye, and L. Bushnell. Stability analysis of netowrked control systems. In *Proceedings of American Control Conference*, pages 2876–2880, 1999.

[127] P. Wan and M.D. Lemmon. Distributed network utility maximization using event-triggered augmented lagrangian methods. In *American Control Conference*, 2009.

[128] P. Wan and M.D. Lemmon. Distributed network utility maximization using event-triggered barrier methods. In *European Control Conference*, Budapest, 2009.

[129] P. Wan and M.D. Lemmon. Event-triggered distributed optimization in sensor networks. In *Information Processing in Sensor Networks (IPSN)2*, San Francisco, CA, USA, 2009.

[130] P. Wan and M.D. Lemmon. An event-triggered distributed primal-dual algorithm for network utility maximization. In *IEEE Conference on Decision and Control*, Shanghai, China, 2009.

[131] P. Wan and M.D. Lemmon. Optimal power flow in microgrids using event-triggered optimization. In *American Control Conference*, Baltimore, Maryland, USA, 2010.

[132] D. Wang, V. Chandar, S.-Y. Chung, and G. Wornell. Error Exponents in Asynchronous Communication. In *Proc. IEEE Int. Symp. Information Theory (ISIT)*, pages 1071–1075, August 2011.

[133] X. Wang and M. Lemmon. Event-triggering in distributed networked control systems. *Automatic Control, IEEE Transactions on*, 56(3):586–601, 2011.

[134] X. Wang and M. Lemmon. On event design in event-triggered feedback systems. *Automatica*, 2011.

[135] X. Wang and M.D. Lemmon. Event-triggering in distributed networked systems with data dropouts and delays. In *Hybrid Systems: computation and control*, San Francisco, CA, USA, 2009.

[136] X. Wang and M.D. Lemmon. Self-triggered feedback control systems with finite-gain L2 stability. *IEEE Transactions on Automatic Control*, 54(3):452–467, 2009.

[137] X. Wang and M.D. Lemmon. Asymptotic staiblity in distributed event-triggered networked control systems with delays. In *Proceedings of the American Control Conference*, 2010.

[138] X. Wang and M.D. Lemmon. Self-triggering under state-independent disturbances. *IEEE Transactions on Automatic Control*, 55(6):1494–1500, 2010.

[139] Xiaofeng Wang and M.D. Lemmon. Attentively efficient controllers for event-triggered feedback systems. In *IEEE Conference on Decision and Control*, 2011.

[140] Q Wen, R. Kumar, J. Huang, and H. Liu. A framework for fault-tolerant control of discrete event systems. *Automatic Control, IEEE Transactions on*, 53(8):1839 –1849, 2008.

[141] A. Willig. Recent and Emerging Topics in Wireless Industrial Communications: A Selection. *IEEE Trans. Industrial Informatics*, 4(2):102–124, May 2008.

[142] A. Willig, K. Matheu, and A. Wolisz. Wireless Technology in Industrial Networks. *Proc. IEEE*, 93(6):1130–1151, June 2005.

[143] A.F.T. Winfield, C.J. Harper, and J. Nembrini. Towards the application of swarm intelligence in safety critical systems. In *System Safety, 2006. The 1st Institution of Engineering and Technology International Conference on*, pages 89–95. IET, 2006.

[144] Tichakorn Wongpiromsarn, Ufuk Topcu, and Richard M. Murray. Receding horizon control for temporal logic specifications. In *Proceedings of the 13th ACM international conference on Hybrid systems: computation and control*, HSCC '10, pages 101–110, New York, NY, USA, 2010. ACM.

[145] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *Networking, IEEE/ACM Transactions on*, 1(3):372–385, 1993.

[146] J. Yi, C. Poellabauer, X.S. Hu, and L. Zhang. Minimum bandwidth reservations for periodic streams in wireless real-time systems. to appear in IEEE Transactions on Mobile Computing, 2010.

[147] Abdellatif Zaidi, Shivaprasad Kotagiri, , J. Nicholas Laneman, and Luc Vandendorpe. Multiaccess Channels with State Known to One Encoder: Another Case of Degraded Message Sets. In *Proc. IEEE Int. Symp. Information Theory (ISIT)*, Seoul, Korea, July 2009.

[148] Abdellatif Zaidi, Shivaprasad Kotagiri, J. Nicholas Laneman, and Luc Vandendorpe. Cooperative Relaying with State Available at the Relay Only. In *Proc. IEEE Information Theory Workshop (ITW)*, Porto, Portugal, May 2008.

[149] Abdellatif Zaidi, Shivaprasad Kotagiri, J. Nicholas Laneman, and Luc Vandendorpe. Cooperative Relaying with State Available Non-Causally at the Relay. *IEEE Trans. Inform. Theory*, 56(5):2272–2298, May 2010.

[150] W. Zhang, M.S. Branicky, and S.M. Phillips. Stability of networked control systems. *IEEE Control Systems Magazine*, 21:84–99, 2001.

[151] Q. Zhu, C. Rieger, and T. Basar. A hierarchical security architecture for cyber-physical systems. In *Resilient Control Systems (ISRCS), 2011 4th International Symposium on*, pages 15–20. IEEE, 2011.