

Taming Uncertainties in Real-Time Routing for Wireless Networked Sensing and Control

Xiaohui Liu, Hongwei Zhang, Qiao Xiang, Xin Che, Xi Ju
Department of Computer Science, Wayne State University, USA
{xiaohui,hongwei,xiangq27,chexin,xiju}@wayne.edu

ABSTRACT

Real-time routing is a basic element of closed-loop, real-time sensing and control, but it is challenging due to dynamic, uncertain link/path delays. The probabilistic nature of link/path delays makes the basic problem of computing the probabilistic distribution of path delays NP-hard, yet quantifying probabilistic path delays is a basic element of real-time routing and may well have to be executed by resource-constrained devices in a distributed manner; the highly-varying nature of link/path delays makes it necessary to adapt to in-situ delay conditions in real-time routing, but it has been observed that delay-based routing can lead to instability, estimation error, and low data delivery performance in general. To address these challenges, we propose the *Multi-Timescale Estimation (MTE)* method; by accurately estimating the mean and variance of per-packet transmission time and by adapting to fast-varying queueing in an accurate, agile manner, MTE enables accurate, agile, and efficient estimation of probabilistic path delay bounds in a distributed manner. Based on MTE, we propose the *Multi-Timescale Adaptation (MTA)* routing protocol; MTA integrates the stability of an ETX-based directed-acyclic-graph (DAG) with the agility of spatiotemporal data flow control within the DAG to ensure real-time data delivery in the presence of dynamics and uncertainties. We also address the challenges of implementing MTE and MTA in resource-constrained devices such as TelosB motes. We evaluate the performance of MTA using the NetEye and Indriya sensor network testbeds. We find that MTA significantly outperforms existing protocols, e.g., improving deadline success ratio by 89% and reducing transmission cost by a factor of 9.7.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols—Routing Protocols

General Terms

Algorithms, Measurement, Performance

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiHoc'12, June 11–14, 2012, Hilton Head Island, SC, USA.
Copyright 2012 ACM 978-1-4503-1281-3/12/06 ...\$10.00.

Keywords

Wireless sensing and control networks, real-time routing, delay quantile estimation, multi-timescale estimation, multi-timescale adaptation

1. INTRODUCTION

Besides deployments for open-loop sensing such as environmental monitoring, embedded wireless networks are increasingly being explored for real-time, closed-loop sensing and control. For instance, wireless networking standards such as the IEEE 802.15.4e, WirelessHART, and ISA SP100.11a have been defined for industrial monitoring and control [3, 43, 39], wireless sensor networks have been deployed for industrial automation [30, 42], and the automotive industry has also been exploring the application of wireless networks to vehicle sensing and control [17, 40]. In wireless networked sensing and control, message passing (or messaging for short) across wireless networks is a basic enabler for coordination among distributed sensors, controllers, and actuators. In supporting mission-critical tasks such as industrial process control, wireless messaging is required to be reliable (i.e., having high delivery ratio) and in real-time [36]. This is because packet loss and large delay usually reduce the system stability (e.g., in proportional-integral control), lengthen the settling time, and increase the maximum overshoot in control [23].

In multi-hop wireless networks, a basis for reliable, real-time messaging is real-time routing which routes data packets from their sources to destinations within specified deadlines. Nonetheless, link/path delays (i.e., the time taken to successfully deliver a packet across a link or a path) are dynamic, uncertain in wireless sensing and control (WSC) networks due to factors such as the spatiotemporal wireless link dynamics and the queueing dynamics along links/paths. The dynamics and uncertainties in link/path delays introduce fundamental *challenges* to real-time routing:

- Firstly, the dynamics and uncertainties make link/path delays probabilistic in nature. Given the delay distributions of individual links along a path, the basic problem of computing the probabilistically guaranteed path delays is NP-hard [21]; in real-time routing where nodes need to identify paths that ensure certain delay bounds, however, quantifying probabilistic path delays is a basic task, and it may well have to be executed by resource-constrained nodes in a distributed manner.
- Secondly, given that link/path delay is a highly varying metric and that it can change at a short timescale of each packet transmission, it is important to adapt to in-situ delay conditions in routing. Yet the highly-varying nature of link/path delays makes it difficult to accurately estimate path delays in a distributed and agile manner, and it has been observed that

delay-adaptive routing can lead to routing instability and low data delivery performance in general [14, 49].

Despite much work in throughput- or energy-efficiency-oriented wireless routing [6, 13, 20, 47], real-time routing is much less studied. Moreover, the existing work that do consider data delivery delay in wireless routing either only try to minimize average path delay without ensuring probabilistic delay bounds [44, 19, 22, 31, 45], or they do not address the challenges that delay uncertainties pose to the task of quantifying probabilistic path delays and the task of addressing instability of delay-adaptive routing [25, 37]. Therefore, how to enable real-time routing in the presence of dynamic, uncertain link/path delays remains an important open problem for real-time wireless networked sensing and control.

Towards enabling routing with probabilistic delay bounds in WSC networks, we propose the *Multi-Timescale Adaptation (MTA)* routing protocol that addresses the aforementioned challenges of dynamic, uncertain link/path delays in real-time routing. In MTA, nodes leverage the different timescales of dynamics to accurately estimate probabilistic path delay bounds in an agile manner and to adapt spatiotemporal data flow control at the same timescales of the dynamics themselves. More specifically, we make the following contributions:

- For accurate, agile estimation of probabilistic path delay bounds, we decompose contributors to path delay uncertainties into two factors: dynamic per-packet transmission time (which we refer to as *packet-time* hereafter) and dynamic queueing along paths. Through detailed experimental analysis, we find that, given a network condition, the distribution of packet-time is quite stable despite the quick variation of instantaneous packet-time. This enables each node to accurately estimate the mean and variance of packet-time from itself to the next-hop along a path. We also observe that the packet-time for different packet transmissions, whether from the same node or from different nodes, are uncorrelated; this enables each node to compute, for a given time instant, the variance of the path delay from itself to the destination node as the sum of the variances of the packet-times for all the packets queued along the path at that instant. Based on these observations, we develop a multi-timescale approach, denoted by *multi-timescale estimation (MTE)*, to accurately estimate the highly-varying mean and variance of path delay by accurately estimating the mean and variance of packet-time and by adapting to fast-varying queueing in an accurate, agile manner. Using the mean and variance of path delay, we evaluate different methods of upper-bounding quantiles, and we identify Chebyshev Inequality as an effective basis for computing probabilistic delay bounds in constant time.
- For enabling adaptivity while addressing instability and low-performance in real-time routing, we propose the *Multi-Timescale Adaptation (MTA)* routing protocol: to facilitate the aforementioned multi-timescale estimation (MTE) and to avoid detrimental instability while ensuring data delivery performance during adaptation, a directed-acyclic-graph (DAG) is maintained at lower frequencies based on the relatively slow-varying link property ETX (i.e., expected number of transmissions taken to successfully deliver a packet), which reflects network throughput, data delivery reliability, and the overall trend of data delivery delay [13, 15]; at higher frequencies and based on the MTE method, the data flow within the DAG is controlled on a per-packet basis to minimize ETX and to ensure packet delivery within the required

probabilistic delay bound. By ensuring overall stability and performance while addressing short-term dynamics at the same time, MTA enables efficient, real-time routing in the presence of complex dynamics and uncertainties.

- We implement MTE and MTA in TinyOS, and we address the challenges of limited memory, limited CPU capability, and the lack of real-time operation support in TinyOS. Besides the running MTA protocol, these implementation strategies may well be of interest to real-time routing in general.
- We evaluate the performance of MTA and other related work in the high-fidelity sensor network testbeds NetEye [4] and Indriya [2]. We find that MTA significantly outperforms existing protocols, e.g., improving deadline success ratio by 89% and reducing transmission cost by a factor of 9.7.

The rest of the paper is organized as follows. We briefly introduce the NetEye and the Indriya testbeds in Section 2. We present the MTE method and the MTA protocol in Sections 3 and 4 respectively, and then we present the measurement study in Section 5. We discuss related work in Section 6, and we make concluding remarks in Section 7.

2. PRELIMINARIES

Our study leverages two publicly available wireless sensor network testbeds NetEye [4] and Indriya [2]. In what follows, we briefly introduce the two testbeds and the traffic patterns we study.

NetEye testbed. NetEye [4] is deployed in a large lab space at Wayne State University. We use a subset of a 15×7 grid of TelosB motes in NetEye, where every two closest neighboring motes are separated by 2 feet. The subset of the grid forms a random network, and it is generated by removing each mote of the 15×7 grid with probability 0.2.

Each of these TelosB motes is equipped with a 3dB signal attenuator and a 2.45GHz monopole antenna. In our measurement study, we set the radio transmission power to be -25dBm (i.e., power level 3 in TinyOS) such that multihop networks can be created and the link reliability is over 90% for links up to 6 feet long. Given the high availability and high fidelity of NetEye, we mainly use NetEye in our measurement study, but we verify key observations using the Indriya testbed too.

Indriya testbed. Indriya [2] is deployed at three floors of the School of Computing at the National University of Singapore. Our measurement study uses all of its 127 TelosB motes, and we use a transmission power of -10dBm (i.e., power level 11 in TinyOS) to generate a well-connected multi-hop network where the link reliability is over 90% for links up to 20 feet long.

Traffic pattern. Using the two testbeds, we study both periodic and event traffic patterns.

For periodic traffic, we study three types of them based on the amount of queueing they introduce when the default TinyOS routing protocol CTP [20] is used: 1) light traffic: no queueing in network; 2) medium traffic: moderate queueing in network but with very rare queue overflow; 3) heavy traffic: severe queueing in network and with frequent queue overflow. To this end, we select one node as the sink and another 10 nodes as traffic sources; the sink and the sources are nearly at opposite positions in both testbeds to create as many routing hops as possible. In NetEye, more specifically, mote 15 is the sink, and motes 61, 62, 63, 64, 76, 77, 79, 91, 92, and 93 are the sources, with each source generating a packet every 1,000ms, 400ms, and 75ms for light, medium, and heavy traffic respectively. To verify observations from NetEye and given that

the medium traffic is the common case in WSC networks, we also study medium traffic scenario in Indriya as follows: mote 105 at the third floor is the sink, and motes 1 . . . 10 at the first floor are the sources, with each source generating a packet every 600ms.

For event traffic, we use a publicly available event traffic trace for a field sensor network deployment [46, 1] to evaluate the performance of different protocols. The traffic trace corresponds to the packets generated in the 7×7 grid of a field mote network when a vehicle passes across the middle of the network. In NetEye, the 7×7 subgrid in the trace data corresponds to the 49 motes that are the farthest from the sink mote 15. In Indriya, the 7×7 subgrid is mapped to motes 1 . . . 49.

3. MULTI-TIMESCALE ESTIMATION OF PATH DELAYS

To enable routing with probabilistic delay guarantees, a basic task is to quantify probabilistic delays along paths. Given a path $P = v_0, v_1, \dots, v_n, v_{n+1}$ from v_0 to v_{n+1} as shown in Figure 1 and assuming that the number of packets queued at node $v_i (i = 0 \dots n)$ is $m_i(t)$ at time t and that the packet transmission scheduling algorithm at each node is first-come-first-serve (FCFS), the instantaneous delay along path P at time t , denoted by $d_P(t)$, is the delay that a packet arriving at v_0 at time t experiences in reaching the destination node v_{n+1} assuming that the number of queued packets at each node does not change while the packet is being delivered to v_{n+1} . Therefore,

$$d_P(t) = \sum_{i=0}^n \sum_{j=1}^{m_i(t)+1} d_i^j(t), \quad (1)$$

where $d_i^j(t)$ is the time taken for v_i to deliver its j -th queued packet to v_{i+1} at time t .

Challenges of highly-varying delay distribution. Given the distributions of $d_i^j(t) (i = 0 \dots n, j = 1 \dots m_i(t) + 1)$, it is NP-hard to compute the distribution of $d_P(t)$ [21]. Since nodes of WSC networks tend to be resource-constrained (e.g., in CPU and memory), it is infeasible to directly compute the distribution of $d_P(t)$ in general. One alternative is to first sample $d_P(t)$ and then use non-parametric approaches (e.g., the P^2 algorithm) to estimate delay quantiles based on these delay samples [27, 34]. Nonetheless, non-parametric quantile estimation usually converges slowly, for instance, taking more than 200 samples [33]. Yet the distribution of path delays varies at a much shorter timescale than the sample size required for non-parametric quantile estimation to converge. As can be seen from (1), in particular, the distribution of path delays vary with the queueing levels (i.e., $m_i(t)$) along paths, and network queueing can vary quickly over time.¹ For the NetEye medium traffic scenario, for instance, Figure 2 show the histogram of the coherence window size for node queueing, where each coherence window of a time series is a maximal consecutive segment of the time series where all the sample values are the same. We see that the coherence window size is two orders of magnitude less than the sample size required for non-parametric quantile estimation to converge. Therefore, the highly-varying nature of path delay distribution makes non-parametric quantile estimation unable to accurately estimate instantaneous path delay quantiles in an agile manner; that is, the distribution of the path delay changes before the estimation converges to an accurate value.

¹As we will discuss shortly, the distribution of packet-time (i.e., $d_i^j(t)$) is quite stable over time, and we leverage the stability of packet-time distribution in designing our multi-timescale estimation (MTE) method.

To circumvent the computational complexity and the inability of accurate, agile estimation of exact path delay quantiles, we propose to identify upper bounds on probabilistic path delays and to use the delay bounds in identifying paths for real-time data delivery. As we will discuss later in this section, upper bounds on path delays can be derived using probability inequalities such as Chebyshev Inequality [38], and the properly identified delay bounds are still orders of magnitude less than the maximum delays, thus enabling effective utilization of network real-time capacity. Most probability inequalities use the mean and/or the standard deviation of the corresponding random variable, thus we need an accurate, agile mechanism of estimating the mean and standard deviation of path delays. Nonetheless, the sample size required for accurate estimation of path delay statistics tend to be quite large, for instance, being greater than 100 in most cases and can be up to 594 [33], which are significantly greater than the coherence window size of node/path queueing levels. In addition, node queueing levels tend to have low autocorrelation, not to mention staying unchanged, for time lags greater than 100 samples, as shown in Figures 3. Thus, the large sample size requirement and the highly-varying nature of path delay distribution makes it impossible to accurately estimate instantaneous mean path delay in an agile manner; that is, the mean path delay changes before the estimation converges to an accurate value.

Multi-Timescale Estimation (MTE). Towards addressing the challenges of highly-varying distribution and statistics of path delay, we decompose contributors to path delay variations into two factors: dynamic packet-time and dynamic queueing. By leveraging the different timescales at which packet-time and queueing vary, we propose the multi-timescale-estimation (MTE) method that accurately estimates the highly-varying mean and variance of path delay by 1) accurately estimating the mean and variance of packet-time at a longer timescale and 2) by adapting to fast-varying queueing at a shorter timescale. In what follows, we elaborate on the design of the MTE method.

OBSERVATION 1. *Packet-time distribution is stable.*

Through detailed experimental analysis, we find that, given a network condition, the distribution of packet-time is quite stable despite the quick variation of instantaneous packet-time. Using the Generalized KPSS test [24], we analyze the stationarity window size of packet-time, where each stationary window of a time series is a maximal consecutive segment of the time series that is weak-sense stationary (i.e., with a constant mean and variance over time). Figure 4 shows the histogram of the stationarity window size for packet-time in the NetEye medium traffic scenario. The minimum and the maximum window size are 1,901 and 21,937 respectively, both of which are significantly greater than the sample size required to precisely estimate the mean and variance of packet-time. For instance, Figure 5 shows the empirical CDF of the sample size required for estimating the mean packet-time at 90% accuracy and 90% confidence level. The sample size requirement is less than 360 with over 99% probability, and the maximum sample size requirement is 600.

The stability of the packet-time distribution has two important implications:

- Given the real-time requirement on data delivery and the constrained memory size in embedded WSC networks, the stationarity window size tends to be greater than the maximum node queue size in general. This implies that, in Formula (1), the mean and variance for $d_i^j(t)$'s ($j = 1 \dots m_i(t) + 1$) are the same for a given $i (i = 0 \dots n)$, and we denote them as




Figure 1: An example path P




Figure 2: Histogram of the coherence window size for node queuing level




Figure 3: Autocorrelation of node queuing levels




Figure 4: Histogram of stationarity window size for packet-time




Figure 5: Empirical CDF of the sample size required for estimating the mean packet-time at 90% accuracy and 90% confidence level




Figure 6: Autocorrelation coefficient of packet-time along the same link




Figure 7: Correlation coefficient between the packet-time across different links along a path

$\mu(d_i(t))$ and $\sigma^2(d_i(t))$. Therefore, the mean path delay can be computed as follows:

$$\mu(d_P(t)) = \sum_{i=0}^n (m_i(t) + 1)\mu(d_i(t)). \quad (2)$$

- The stability of packet-time distribution enables each node to accurately estimate the mean and variance of packet-time from itself to the next-hop along a path, and, with Formula (2), the mean path delay can be accurately estimated.

OBSERVATION 2. *Packet-time is uncorrelated.*

We also observe that the packet-time for different transmissions, whether from the same node or from different nodes, tend to be uncorrelated. For instance, Figure 6 shows the small autocorrelation coefficient of packet-time along the same link, with the boxplot showing, for each lag, the distribution of autocorrelation coefficients across different links; Figure 7 shows the small correlation coefficient of packet-time along different links, where the lag is defined as the number of packets queued between two packets at different node queues as shown in Figure 1, assuming that the packets queued at time t flow through the nodes queues in a FIFO manner. We also observe that the median autocorrelation coefficient and the median cross-link correlation coefficient of packet-time is zero at 99% confidence level. An intuition for the uncorrelatedness between packet-time is that, given a network condition, the behavior of one packet transmission does not have much impact on the behavior of another packet transmission as far as the MAC protocol is concerned.

The Bienaymé Formula [35] shows that the variance of the sum of pair-wise uncorrelated random variables is the sum of the variances of the individual random variables. Thus our observations on the uncorrelatedness of packet-time along a path enables each

node to compute, for a given time instant, the standard deviation of the path delay from itself to the destination node as the square root of the sum of the variances of the packet-times for all the packets queued along the path at that instant. For path P of Figure 1, for instance, the standard deviation of $d_P(t)$ can be computed as follows:

$$\sigma(d_P(t)) = \sqrt{\sum_{i=0}^n (m_i(t) + 1)\sigma^2(d_i(t))}. \quad (3)$$

For a typical 5-hop path in NetEye, Figure 8 shows the histogram of the relative errors in estimating the standard deviation of path delay in the presence of different queuing levels along the path. We see that the estimation is quite accurate, with most relative errors within the range of $(-0.075, 0.075)$. Note that, if we directly estimate the variance of the sojourn time $d_i^{so}(t) = \sum_{j=1}^{m_i(t)+1} d_i^j(t)$ at node v_i without decomposing $d_i^{so}(t)$ into its individual components $d_i^j(t)$'s ($j = 1 \dots m_i(t) + 1$), we cannot compute $\sigma(d_P(t))$ as $\sqrt{\sum_{i=0}^n \sigma^2(d_i^{so}(t))}$. This is because, for $i1 \neq i2$, $d_{i1}^{so}(t)$ and $d_{i2}^{so}(t)$ are correlated due to the correlation between queuing levels at different nodes of a path.

In the simplified scenario of Figure 1, the next-hops of all the packets in a node queue are the same. In reality, a node may well use different next-hops for different packets, for instance, depending on packet deadlines as we will discuss in Section 4. Assume that, at time t , each node v_i ($i = 0 \dots n$) has $N_i(t)$ number of next-hops, the number of packets (including the one arriving at v_0 at time t) to be forwarded to its k -th next hop is $m_i^k(t)$, and the packet-time from v_i to its k -th next-hop is $d_{i,k}(t)$. We observe that, given v_i and t , $d_{i,k}(t)$'s tend to be uncorrelated for different k 's. For instance, Figure 9 shows the histogram of the correlation coefficient between the packet-time across different outgoing links from the same node. The correlation coefficient is very small and




Figure 8: Histogram of relative errors in estimating the standard deviation of path delay




Figure 9: Histogram of the correlation coefficient between the packet-time across different outgoing links from the same node




Figure 10: CDF of link queueing level changes




Figure 11: Bounds on 90th percentile of path delay (more readable in color)

is less than 3% most of the time. We also observe that the median correlation coefficient is zero at 99% confidence level. Then, based on Formulae (2) and (3), the mean and standard deviation of the delay along a path P can be computed as follows:

$$\begin{aligned} \mu(d_P(t)) &= \frac{\sum_{i=0}^n \sum_{k=1}^{N_i(t)} m_i^k(t) \mu(d_{i,k}(t))}{\sum_{i=0}^n \sum_{k=1}^{N_i(t)} m_i^k(t)} \\ \sigma(d_P(t)) &= \sqrt{\sum_{i=0}^n \sum_{k=1}^{N_i(t)} m_i^k(t) \sigma^2(d_{i,k}(t))} \end{aligned} \quad (4)$$

OBSERVATION 3. *Network queueing is relatively stable at short timescales.*

To leverage Formula (4) in estimating the mean and variance of the delay along path P in a distributed manner, each node v_i can compute the mean and variance of the delay from itself to the destination node v_{n+1} based on those of its next-hop along P . Denoting the mean and standard deviation of the delay from v_i to v_{n+1} by $\mu(d_P^i(t))$ and $\sigma(d_P^i(t))$ respectively ($i = 0 \dots n$), then we have

$$\begin{aligned} \mu(d_P^i(t)) &= \mu(d_P^{i+1}(t)) + \sum_{k=1}^{N_i(t)} m_i^k(t) \mu(d_{i,k}(t)) \\ \mu(d_P^{n+1}(t)) &= 0 \end{aligned} \quad (5)$$

$$\begin{aligned} \sigma(d_P^i(t)) &= \sqrt{\sigma^2(d_P^{i+1}(t)) + \sum_{k=1}^{N_i(t)} m_i^k(t) \sigma^2(d_{i,k}(t))} \\ \sigma(d_P^{n+1}(t)) &= 0 \end{aligned} \quad (6)$$

Formulae (7) and (8) can be implemented using a distance-vector-type routing algorithm. Due to information diffusion delay τ_i from v_{i+1} to v_i in routing, the implemented version of Formulae (7) and (8) are

$$\begin{aligned} \mu(d_P^i(t)) &= \mu(d_P^{i+1}(t - \tau_i)) + \sum_{k=1}^{N_i(t)} m_i^k(t) \mu(d_{i,k}(t)) \\ \mu(d_P^{n+1}(t)) &= 0 \end{aligned} \quad (7)$$

$$\begin{aligned} \sigma(d_P^i(t)) &= \sqrt{\sigma^2(d_P^{i+1}(t - \tau_i)) + \sum_{k=1}^{N_i(t)} m_i^k(t) \sigma^2(d_{i,k}(t))} \\ \sigma(d_P^{n+1}(t)) &= 0 \end{aligned} \quad (8)$$

For accurate estimation of $\mu(d_P^i(t))$ and $\sigma^2(d_P^i(t))$, we need to make τ_i as small as possible. This can be achieved by having v_{i+1} piggyback $\mu(d_P^{i+1}(t))$ and $\sigma^2(d_P^{i+1}(t))$ onto its data transmissions as well as control signaling (e.g., broadcast of routing beacons) so that v_i can overhear the values quickly. In WSC networks, sample data about physical behavior are usually generated periodically in a continuous manner, thus τ_i is at the same timescale of inter-packet arrival interval, which enables quick diffusion of path delay statistics.

We also observe that, even though network queueing varies significantly at a long timescale of hundreds of inter-packet intervals,

it is much more stable at a short timescale of a few inter-packet intervals. In the NetEye medium traffic scenario and for the time lags of 1, 5, and 10 packet transmissions, for instance, Figure 10 shows the empirical cumulative distribution function (CDF) of link queueing level changes (i.e., changes in $m_i^k(\cdot)$). We see that, at the timescale of a few inter-packet intervals, network queueing remains relatively stable, and, with more than 90% probability, the absolute changes in link queueing levels are no more than 1. To ensure enough real-time capacity for each source node, routing hops tend to be limited (e.g., less than 10) in WSC networks [5]. This fact, together with the quick information diffusion and the relative stable network queueing at short timescales, enables the MTE method to accurately estimate the mean and variance of path delays in an agile manner.

Probabilistic path delay bound. With the mean and variance of path delay estimated via the MTE method, a node can derive the probabilistic path delay bounds using probability inequalities. To this end, we have

PROPOSITION 1. *For a random variable X , if*

$$Pr\{X \geq f(x)\} \leq g(x), \quad (9)$$

then $Q_X^q = f(g^{-1}(1-q))$ is an upper bound on the q -quantile of X , where g^{-1} is the inverse function of g .

PROOF. Let $g(x) = 1 - q$, then

$$Pr\{X \geq f(x)\} \leq 1 - q.$$

Thus,

$$Pr\{X \leq f(x)\} \geq q$$

Since $x = g^{-1}(1 - q)$, we have

$$Pr\{X \leq f(g^{-1}(1 - q))\} \geq q$$

Thus, $f(g^{-1}(1 - q))$ is an upper bound on the q -quantile of X . \square

For $\lambda > 0$ and a non-negative random variable X with mean μ and variance σ^2 , two widely-applicable probability inequalities [41] are the Markov Inequality

$$Pr\{X \geq \lambda\} \leq \frac{\mu(X)}{\lambda}, \quad (10)$$

and the one-tailed Chebyshev Inequality

$$Pr\{X - \mu \geq \lambda\sigma\} \leq \frac{1}{1 + \lambda^2}. \quad (11)$$

Thus, we have

COROLLARY 1. *Using Markov Inequality, $Q_X^q = \frac{\mu}{1-q}$.*

COROLLARY 2. *Using one-tailed Chebyshev Inequality, $Q_X^q = \mu + \sigma \sqrt{\frac{q}{1-q}}$.*

For a typical five-hop path in NetEye, Figure 11 shows the ground truth and the estimated 90-percentile of the path delay when it is the sum of 5 and 40 packet-time random variables respectively. For comparison purpose, we present the probabilistic delay bounds estimated via the optimal-partition-minimum-delay (OPMD) method [34] or by assuming path delay is normally distributed [29]; we also present the maximum path delay. Given that Markov Inequality usually gives a looser bound than one-tailed Chebyshev Inequality, the probabilistic delay bound by Markov Inequality is greater than that by Chebyshev Inequality. Compared with the method of assuming path delay is normally distributed, the bound by Chebyshev Inequality is always greater than the actual 90-percentile path delay, whereas the former method underestimates the 90-percentile for the case of 40 packet-time random variables. Given that the OPMD method is rather conservative in estimating path delay bound, the bound by Chebyshev Inequality is also much less than the bound by the OPMD method, especially when path queueing increases. We also see that the maximum path delay is orders of magnitude greater than the bound by Chebyshev Inequality, thus using probabilistic delay bound instead of maximum path delay helps improve application-usable real-time capacity. Since the bound by Chebyshev Inequality upper-bounds and is close to the probabilistic path delay, we use the probabilistic delay bound by Chebyshev Inequality in our protocol design.

From FCFS to EDF. In real-time scheduling, the earliest-deadline-first (EDF) algorithm is a commonly used algorithm, and it can ensure a smaller deadline than what is feasible with the first-come-first-serve (FCFS) algorithm. Our MTE method of computing probabilistic path delay bounds is derived by assuming that nodes use FCFS algorithm in intra-node transmission scheduling, but we observe that the delay bounds derived via our FCFS-based MTE method can also serve as the basis of selecting real-time packet forwarding paths even if the EDF algorithm is used. Formally,

PROPOSITION 2. *If routed along a path whose probabilistic delay bound computed via the FCFS-based MTE method is less than the relative deadline of a packet, the packet will reach its destination before its deadline even if the EDF algorithm is used for intra-node transmission scheduling.*

(Interested readers can find the proof in [33].)

4. MULTI-TIMESCALE ADAPTATION FOR REAL-TIME ROUTING

Real-time routing is subject to dynamics and uncertainties at multiple timescales. At a longer timescale, link properties such as ETX (i.e., expected-number-transmissions to successfully deliver a packet) vary as a result of changing environmental conditions (e.g., temperature); at a shorter timescale, data transmission delay varies on a per-packet basis, and bursty traffic may introduce sudden changes to network conditions. Robust system design usually requires adaptation to dynamics at the same or shorter timescales of the dynamics themselves. Yet we have found that, due to the highly-varying nature of path delays, routing using delay-based metrics can introduce large estimation errors and lead to routing instability as well as low performance [49]. To ensure long-term stability and optimality while addressing short-term dynamics at the

same time, we propose a multi-timescale adaptation (MTA) framework for real-time routing as follows.

At lower frequencies, a directed-acyclic-graph (DAG) is maintained for data forwarding, and any path within the DAG is a candidate path for packet delivery. Given that link/path ETX reflects network throughput, data delivery reliability, and the overall trend of data delivery delay [13, 15], and that ETX-based routing structures tend to be stable even if ETX is dynamic [48], we propose to maintain the data forwarding DAG based on link and path ETX such that the DAG reflects long-term system optimality and changes relatively slowly compared with delay variation. More specifically, there is a directed edge from node v_i to v'_i in the DAG if and only if the minimum path ETX from v'_i to the destination v'_0 is less than that from v_i to v'_0 . The DAG defines, for each node v_i , a set of forwarder candidates $\mathcal{R}(v_i)$ where $v'_i \in \mathcal{R}(v_i)$ if and only if link $\langle v_i, v'_i \rangle$ belongs to the DAG.

At higher frequencies, the spatiotemporal flow of packets within the data forwarding DAG is adaptively controlled to ensure reliable, real-time data delivery in the presence of short-timescale dynamics such as transient packet losses and per-packet variations of link delay. More specifically, each packet contains information about the remaining time to deadline, denoted by L , and the required real-time guarantee probability q . When the packet reaches a node v_i , v_i first finds the set of forwarder candidates within the DAG, denoted by $\mathcal{R}'(v_i, L, q)$, that can ensure the real-time requirements L and q ; then v_i sets the next-hop node of the packet as the node of the smallest path ETX to v'_0 among all the nodes in $\mathcal{R}'(v_i, L, q)$, and v_i puts the packet in the transmission queue. Queued packets are then scheduled for transmission using the earliest-deadline-first (EDF) algorithm.

Using the above approach, packets are always routed along the minimum-ETX path that satisfies the real-time data delivery requirement. In the presence of heavy traffic load that induces queueing, this real-time forwarding mechanism creates the water-filling effect as follows: packets are delivered to the minimum-ETX path until the path cannot ensure the required timeliness of data delivery (e.g., due to queueing), at which point the forthcoming packets are delivered to the path with the second-minimum-ETX, and so on; once the delay along paths of less ETX decreases (e.g., due to reduced queueing), more packets will be delivered to those paths to fill them up. This process repeats at the same timescale of packet arrival process, and it enables real-time data delivery while ensuring as small ETX in data delivery as possible. Unlike traditional delay-based routing that may lead to instability and low performance, this quick adaptation of spatial packet flow is enabled by 1) the MTE method of accurate, agile estimation of probabilistic path delay bound at the same timescale of the changes in path delay distribution and by 2) the overall stability of spatial packet flow along the data forwarding DAG.

Due to the limitation of space, we delegate the system architecture and the implementation strategies for the MTA framework to [33].

5. MEASUREMENT EVALUATION

5.1 Methodology

Protocols. We have implemented the MTA framework (including the MTE method) in TinyOS. To understand the design decisions of MTA, we have comparatively studied MTA with its variants; due to the limitation of space, we delegate the detailed discussions to [33]. Towards understanding the benefits of MTA, we also comparatively study MTA with the following existing protocols:

- *MCMP*: a multi-path QoS routing protocol where end-to-end QoS requirements on reliability and timeliness are uniformly divided into per-hop reliability and timeliness requirements, upon which a node chooses the minimum number of next-hops to satisfy the per-hop requirements in data delivery [25];
- *MM*: the geographic routing protocol MMSPEED [19] that routes and schedules packet transmissions based on nodes' distances to destinations, packet delivery deadlines, and mean link delays; MMSPEED also tries to improve packet delivery reliability by transmitting packets along multiple paths; (Note: we denote MMSPEED as MM for the readability of figures to be presented in the next subsection.)
- *MM-CD*: same as MMSPEED but, instead of using the mean link delay, uses a conservative estimate of link delay that equals the sum of the mean delay and three times the standard deviation of the delay;
- *SDRCS*: similar to MMSPEED but, instead of using geographic distance, uses data-forwarding hop-count as the measure of distance, where the hop-count is computed based on received-signal-strength (RSS) between nodes [44]; data forwarding is through receiver contention similar to that in opportunistic routing;
- *CTP* : a ETX-based non-real-time routing protocol in TinyOS [20].

(Note: for MM and MM-CD which use node locations, we manually configured each node with the correct location in our TinyOS programs.)

Traffic & real-time requirements. We use all periodic and event traffic discussed in Section 2. In our experiments, we use 90% as the required real-time guarantee probability by default, but we have also experimented with the real-time guarantee probability of 99% and observed similar phenomena [33]. For differentiating the performance of different protocols, the deadline for each traffic scenario is chosen so that it is neither too stringent (that no protocol can support) nor too loose (that all protocols can support). In NetEye, the deadlines for light, medium, and heavy traffic are 250ms, 2 seconds, and 7.5 seconds respectively, and the deadline for event traffic is 10 seconds. In Indriya, the deadlines for periodic and event traffic are 2 seconds and 2.5 seconds respectively.

Due to the limitation of space, here we only present results for medium periodic traffic in NetEye and Indriya. Similar phenomena have been observed for other traffic scenarios; interested readers can find them in [33]. We have also experimented with networks of lower connectivity and have observed similar phenomena [33].

Metrics. For each combination of protocol, testbed, and traffic scenario, we run it for 10 times and evaluate protocol performance in terms of the following metrics:

- *Deadline success ratio (DSR)*: ratio of packets delivered to the sink before their deadlines;
- *Packet delivery ratio (PDR)*: ratio of packets delivered to the sink;
- *Number of transmissions per packet delivered (NTX)*: total number of transmissions, including retransmissions, divided by the number of unique packets delivered to the sink.

To understand protocol behavior in more detail, we also analyze the different causes for a packet to miss its deadline:

- *Overflow*: packet discarded due to node queue overflow;

- *Transmission failure*: a packet not delivered to the next hop even after the maximum number of retransmissions at a node;
- *Rejection*: no candidate path can ensure the required real-time delivery guarantee (i.e., deadline and probability) of a packet;
- *Expiration*: deadline expired before the packet reaches the sink, whether or not the packet is delivered to the sink.

5.2 Measurement results

NetEye. For periodic medium traffic in NetEye, Figures 12, 13, 14, and 15 show the deadline success ratio, packet delivery ratio, transmission cost, and packet delivery status respectively. We see that MTA always ensures the required real-time delivery performance, and the median deadline success ratio in MTA is 76%, 89%, 86%, 11%, and 38% higher than that in MCMP, MM, MM-CD, SDRCS, CTP respectively. The median number of transmissions per packet delivered in MTA is less than that in MCMP, MM, MM-CD, SDRCS, CTP by a factor of 3.7, 9.7, 6.5, 1.2, and 1.2 respectively.

One reason why the deadline success ratio in MCMP, MM, and MM-CD is very low is because a large fraction of packets are lost due to queue overflow as can be seen in Figure 15. They all try to use multiple paths to ensure data delivery reliability: at the sources, multiple copies of a packet can be sent; at the next hop, each of these copies can be multiplied again, and so on. This multi-path routing mechanism can lead to exponentially increasing number of copies of a packet, thus causing severe queue overflow and large transmission cost as shown in Figures 15 and 14 respectively. SDRCS does not have this problem because it does not use multi-path routing, instead it uses a data forwarding mechanism similar to opportunistic routing.

Another reason why MCMP, MM, MM-CD, and SDRCS do not perform well is because, by evenly dividing end-to-end QoS requirements into per-hop requirements, they implicitly assume that network conditions are uniform across the network which is usually not the case. Among these protocols, the negative impact of this assumption is relatively less severe in SDRCS because it uses signal strength as the basis of measuring forwarding distances and signal strength is a better metric for measuring wireless link quality than geographic distance. MTA does not have this problem because MTE enables accurate, agile estimation of end-to-end delay quantiles without assuming uniform network conditions. A third reason for the low performance of MM, MM-CD, and SDRCS is because they only consider mean delays instead of the probabilistic distributions of delays.

Compared with MTA, CTP has higher packet delivery ratio, but CTP only enables a median deadline success ratio of 56% which is lower than the real-time guarantee probability of 90% and much lower than the 93% probability guarantee by MTA. This is because CTP only considers path ETX in routing, and it is delay-unaware. Even if a low-ETX path is experiencing large delay due to queuing, CTP still uses the path, thus leading to large data delivery delay and deadline miss. Through accurate, agile estimation of path delay via MTE, in contrast, MTA can switch to a less congested path whenever it detects the inability of the current low-ETX path to deliver packets before their deadlines, creating the water-filling effect as we have discussed in Section 4. For instance, Figure 5.2 shows the histogram of the ETX of the paths taken by all the packets. We see that MTA tends to use paths of lower-ETX with higher probability, even though the minimum-ETX path will not always be used with the highest probability (e.g., when the capacity of the




Figure 12: Deadline success ratio: MTA and existing protocols in NetEye




Figure 13: Packet delivery ratio: MTA and existing protocols in NetEye




Figure 14: Number of transmissions per packet delivered: MTA and existing protocols in NetEye




Figure 15: Packet delivery status: MTA and existing protocols in NetEye




Figure 17: Deadline success ratio: MTA and existing protocols in Indriya




Figure 18: Packet delivery ratio: MTA and existing protocols in Indriya




Figure 19: Number of transmissions per packet delivered: MTA and existing protocols in Indriya




Figure 20: Packet delivery status: MTA and existing protocols in Indriya




Figure 16: Histogram of the ETX of the paths taken by all the packets

minimum-ETX path is reduced due to the shared path segments with other paths).

Indriya. Figures 17, 18, 19, and 20 show the deadline success ratio, packet delivery ratio, transmission cost, and packet delivery status for periodic traffic in Indriya respectively. The overall relative behavior between protocols is similar to that in NetEye, but the performance of MM, MM-CD, and SDRCS become much worse compared with MTA. MTA still ensures real-time data delivery, but MM and MM-CD can hardly deliver any packet to the sink, let alone delivering packets in time; the deadline success ratio in SDRCS is also more than 72% less than that in MTA. One major cause for this is that MM, MM-CD, and SDRCS implicitly assume uniform network conditions while the degree of heterogeneity in Indriya is significant and higher than that in NetEye. As a result of the uniformity assumption, for instance, about 60% and 50% of packets are rejected in MM and MM-CD respectively as shown in Figure 20. Given a packet, more specifically, its deadline and the distance from its source to the sink determines the required forwarding speed for this packet. For the packet to reach the sink,

every single hop it traverses has to provide a speed no less than the required speed. That is, if any intermediate hop cannot meet the required speed, a packet is rejected. For MM and MM-CD, it only takes about 3 or 4 hops to reach the sink from the sources in NetEye, while it takes about 7 or 8 hops in Indriya, making packets in Indriya more likely to be rejected. Note that there is significant queue overflow in SDRCS because SDRCS happen to use low reliability links in Indriya which reduce the network throughput and thus increase the queueing and queue overflow.

6. RELATED WORK

QoS routing has been well studied for the Internet [8, 11] and wireless networks [18, 9, 10]. But most did not consider uncertainties in link/path properties (e.g., delay). Link property uncertainties were considered in [21] and [29], and it was shown that the problem of checking probabilistically guaranteed path delays is NP-hard [21]. Focusing on Internet QoS routing, these work assumed link-state routing, and their solutions were not amenable to light-weight, distance-vector-type implementation. Since link-state routing is usually not suitable for dynamic, resource constrained WSC networks where reliable network-wide link-state update itself is a challenging issue and where nodes may only have very limited memory space (e.g., up to 4KB of RAM), the approaches of [21] and [29] are not applicable to WSC networks.

Data delivery delay was also considered in wireless and sensor network routing [19, 22, 25, 31, 45]; but they only tried to minimize average path delay without ensuring probabilistic delay bounds [44, 19, 22, 31, 45], they did not consider the probabilistic nature of link/path delays [37], they were based on geographic forwarding without addressing network non-uniformity and wireless communication irregularity [22, 19], or they uniformly partitioned multi-hop QoS requirements (e.g., reliability and timeliness) along the

links of a path without considering network non-uniformity [44, 25]. Huang et al. [25] used Chebyshev Inequality in single-hop real-time satisfiability testing, but they did not address the challenges of accurate, agile estimation of multi-hop probabilistic delay bounds, and they did not comparatively study Chebyshev Inequality with other well-known probability inequalities. Liu et al. [34] proposed the pseudo-polynomial time algorithm optimal-partition-minimum-delay (OPMD) for upper-bounding probabilistic path delays, but, as we have shown in Section 3, the bound of the OPMD algorithm is quite loose for multi-hop paths.

Multi-timescale adaptation has been considered in Internet traffic engineering [7, 28]. Focusing on load balancing, these work did not consider QoS assurance. Liu et al. [34] also studied multi-timescale adaptation in routing, but they used the OPMD method to estimate probabilistic delay bounds which are significantly looser than the bound identified through our MTE method, thus leading to real-time capacity loss. The IETF ROLL working group [26] considered building routing trees based on directed-acyclic-graphs (DAG) for low-power wireless networks. Serving as a general reference framework, the ROLL routing proposal did not consider specific optimization methods (e.g., for real-time guarantees). For stable data delivery reliability, Lin et al. [32] proposed to route data based on long-term link properties and to address transient perturbations using power and retransmission control; they focused on data delivery reliability instead of real-time, thus they did not consider the challenges of dynamic, uncertain link/path delays in real-time routing.

The WirelessHART [43] and the ISA SP100.11a [39] standards have been recently proposed for wireless networking in industrial process measurement and control. They mostly focus on high-level system frameworks instead of specific algorithms in real-time routing. In the literature of real-time wireless networking, techniques such as power control [12] as well as joint routing and scheduling [37] have been studied; energy-efficiency [16] has also been considered too. Orthogonal to these studies, our study here has focused on addressing the challenges that dynamic, uncertain link/path delays pose to two basic elements of real-time routing, i.e., determining probabilistic path delays and addressing instability in delay-adaptive routing. Integrating our results with those work will be an interesting research avenue to pursue, but detailed study of it is beyond the scope of this paper.

7. CONCLUDING REMARKS

For addressing the challenges of highly-varying path delays to distributed estimation of path delay quantiles, we have proposed the MTE method that leverages the stability of packet-time distribution and the quick diffusion of path delay statistics (i.e., mean and variance) to accurately estimate probabilistic path delay bounds in an agile manner. Based on accurate, agile characterization of path delays using MTE, our MTA routing framework enables the stability and optimality of data forwarding while adapting to fast-changing network queueing and delay. Through extensive measurement study in both the NetEye and the Indriya wireless sensor network testbeds, we have shown that MTE/MTA-based routing ensures efficient, real-time data delivery, and it significantly outperforms existing real-time routing protocols. We have mainly focused on real-time spatial flow control in this study, even though we have experimentally analyzed the benefits of using EDF instead of FCFS in intra-node scheduling; how to control temporal packet flow between neighbors and across the network and how to jointly optimize the spatial and temporal packet flow will be an important area to explore, where the MTE method and the MTA framework are expected to serve as basic systems building-blocks. The tech-

nique of leveraging different timescales of dynamics in protocol design may well be of generic interest to wireless networking in dynamic, uncertain environments too.

Acknowledgment

We thank George Yin for discussions on probability inequalities, and we also thank our MobiHoc'12 shepherd Ivan Seskar for helpful comments on our paper. The authors' work is supported in part by NSF awards CNS-1136007, CNS-1054634, GENI-1890, and GENI-1633, as well as grants from Ford Research and GM Research.

8. REFERENCES

- [1] An event traffic trace for sensor networks. <http://www.cs.wayne.edu/~h Zhang/group/publications/Lites-trace.txt>.
- [2] Indriya testbed. <http://indriya.comp.nus.edu.sg/>.
- [3] IEEE 802.15.4e Working Group. <http://www.ieee802.org/15/pub/TG4e.html>.
- [4] NetEye testbed. <http://neteye.cs.wayne.edu/neteye/home.php>, 2008.
- [5] S. Bapat, V. Kulathumani, and A. Arora. Analyzing the yield of exscal, a large-scale wireless sensor network experiment. In *IEEE ICNP*, pages 53–62, 2005.
- [6] S. Biswas and R. Morris. ExOR: Opportunistic multi-hop routing for wireless networks. In *ACM SIGCOMM*, 2005.
- [7] M. Caesar, M. Casado, T. Koponen, J. Rexford, and S. Shenker. Dynamic route computation considered harmful. *ACM SIGCOMM Computer Communication Review*, 40(2), 2010.
- [8] A. Chakrabarti and G. Manimaran. Reliability constrained routing in QoS networks. *IEEE/ACM Transactions on Networking*, 13(3):662–675, 2005.
- [9] S. Chakrabarti and A. Mishra. QoS issues in ad hoc wireless networks. *IEEE Communications Magazine*, pages 142–148, February 2001.
- [10] L. Chen and W. B. Heinzelman. QoS-aware routing based on bandwidth estimation for mobile ad hoc networks. *IEEE JSAC*, 23(3):561–571, 2005.
- [11] S. Chen, M. Song, and S. Sahn. Two techniques for fast computation of constrained shortest paths. *IEEE/ACM Transactions on Networking*, 16(1), February 2008.
- [12] O. Chipara, Z. He, G. Xing, Q. Chen, X. Wang, C. Lu, J. Stankovic, and T. Abdelzaher. Real-time power-aware routing in sensor networks. In *IWQoS*, 2006.
- [13] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *ACM MobiCom*, 2003.
- [14] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. In *ACM SIGCOMM*, 2004.
- [15] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *ACM MobiCom*, 2004.
- [16] S. C. Ergen and P. Varaiya. Energy efficient routing with delay guarantee for sensor networks. *Wireless Networks*, 13(5), 2007.
- [17] Y. P. Fallah, C. Huang, R. Sengupta, and H. Krishnan. Design of cooperative vehicle safety systems based on tight

- coupling of communication, computing and physical vehicle dynamics. In *ACM/IEEE ICCPS*, 2010.
- [18] X. Fang, D. Yang, P. Gundecha, and G. Xue. Multi-constrained anypath routing in wireless mesh networks. In *IEEE SECON*, 2010.
- [19] E. Felemban, C.-G. Lee, E. Ekici, R. Boder, and S. Vural. Probabilistic QoS guarantee in reliability and timeliness domains in wireless sensor networks. In *IEEE INFOCOM*, 2005.
- [20] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *ACM SenSys*, 2009.
- [21] R. A. Guerin and A. Orda. QoS routing in networks with inaccurate information: Theory and algorithms. *IEEE/ACM Transactions on Networking*, 7(3):350–364, 1999.
- [22] T. He, J. Stankovic, C. Lu, and T. Abdelzaher. SPEED: A stateless protocol for real-time communication in sensor networks. In *IEEE ICDCS*, 2003.
- [23] J. Hellerstein, Y. Diao, S. Parekh, and D. M. Tilbury. *Feedback Control of Computing Systems*. Wiley-IEEE Press, 2004.
- [24] B. Hobijn, P. H. Franses, and M. Ooms. Generalizations of the KPSS-test for stationarity. *Statistica Neerlandica*, 58(4):483–502, 2004.
- [25] X. Huang and Y. Fang. Multiconstrained QoS multipath routing in wireless sensor networks. *Wireless Networks*, 14:465–478, 2008.
- [26] IETF. Routing over low power and lossy networks (ROLL) working group. <http://www.ietf.org/html.charters/roll-charter.html>.
- [27] R. Jain and I. Chlamtac. The P^2 algorithm for dynamic calculation of quantiles and histograms without storing observations. *Communications of the ACM*, 28(10), 1985.
- [28] S. Kandula, D. Katabi, B. Davie, and A. Charny. Walking the tightrope: Responsive yet stable traffic engineering. In *ACM SIGCOMM*, 2005.
- [29] T. Korkmaz and M. Krunz. Bandwidth-delay constrained path selection under inaccurate state information. *IEEE/ACM Transactions on Networking*, 11(3), 2003.
- [30] A. LaJoie. Wireless sensors enjoy timely role in factory automation. *Industrial Ethernet Book*, April 2010.
- [31] H. Li, Y. Cheng, and C. Zhou. Minimizing end-to-end delay: A novel routing metric for multi-radio wireless mesh networks. In *IEEE INFOCOM*, 2009.
- [32] S. Lin, G. Zhou, K. Whitehouse, Y. Wu, J. A. Stankovic, and T. He. Towards stable network performance in wireless sensor networks. In *IEEE RTSS*, 2009.
- [33] X. Liu, H. Zhang, and Q. Xiang. Taming uncertainties in real-time routing for wireless networked sensing and control. Technical Report DNC-TR-11-04 (<https://sites.google.com/site/dnctrns/DNC-TR-11-04.pdf>), Wayne State University, 2011.
- [34] X. Liu, H. Zhang, and Q. Xiang. Towards predictable real-time routing for wireless networked sensing and control. In *CPS Week RealWin workshop*, 2011.
- [35] M. Loeve. *Probability Theory*. Springer-Verlag, 1977.
- [36] J. R. Moyne and D. M. Tilbury. Control and communication challenges in networked real-time systems. *Proceedings of the IEEE*, 95(1), 2007.
- [37] S. Munir, S. Lin, E. Hoque, S. M. S. Nirjon, J. A. Stankovic, and K. Whitehouse. Addressing burstiness for reliable communication and latency bound generation in wireless sensor networks. In *IEEE/ACM IPSN*, 2010.
- [38] S. M. Ross. *Introduction to Probability Models*. Academic Press, 2006.
- [39] ISA SP100.11a. <http://www.isa.org/MSTemplate.cfm?MicrositeID=1134&CommitteeID=6891>.
- [40] H.-M. Tsai, C. Saraydar, T. Talty, M. Ames, A. Macdonald, and O. K. Tonguz. Zigbee-based intra-car wireless sensor network. In *IEEE ICC*, 2007.
- [41] J. V. Uspensky. *Introduction to Mathematical Probability*. McGraw-Hill Book Company Inc., 1937.
- [42] A. Willig. Recent and emerging topics in wireless industrial communications: A selection. *IEEE Transactions on Industrial Informatics*, 4(2), May 2008.
- [43] WirelessHART. http://www.hartcomm2.org/hart_protocol/wireless_hart/wireless_hart_main.html.
- [44] Y. Xue, B. Ramamurthy, and M. C. Vuran. SDRCS: A service-differentiated real-time communication scheme for event sensing in wireless sensor networks. *Computer Networks*, 55:3287–3302, 2011.
- [45] S. Yin, Y. Xiong, Q. Zhang, and X. Lin. Traffic-aware routing for real-time communications in wireless multi-hop networks. *Wireless Communications and Mobile Computing*, 6:825–843, 2006.
- [46] H. Zhang, A. Arora, Y. ri Choi, and M. Gouda. Reliable bursty convergecast in wireless sensor networks. *Computer Communications (Elsevier), Special Issue on Sensor-Actuator Networks*, 30(13), 2007.
- [47] H. Zhang, A. Arora, and P. Sinha. Link estimation and routing in sensor network backbones: Beacon-based or data-driven? *IEEE Transactions on Mobile Computing*, 8(5), May 2009.
- [48] H. Zhang, L. Sang, and A. Arora. On biased link sampling in data-driven link estimation and routing in low-power wireless networks. *ACM Transactions on Autonomous and Adaptive Systems*, 4(3), July 2009.
- [49] H. Zhang, L. Sang, and A. Arora. Comparison of data-drive link estimation methods in low-power wireless networks. *IEEE Transactions on Mobile Computing*, 9(11), November 2010.