# Extended Abstract: Real-Time Implications of Multiple Transmission Rates in Wireless Networks

Vartika Bhandari[*]
Dept. of Computer Science
University of Illinois at
Urbana-Champaign
vbhandar@uiuc.edu

Vivek Raghunathan[†]
Dept. of Electrical and
Computer Engineering
University of Illinois at
Urbana-Champaign
vivek@control.csl.uiuc.edu

Bach Duy Bui
Dept. of Computer Science
University of Illinois at
Urbana-Champaign
bachbui2@uiuc.edu

Marco Caccamo[‡]
Dept. of Computer Science
University of Illinois at
Urbana-Champaign
mcaccamo@cs.uiuc.edu

## ABSTRACT

Wireless networks are increasingly being used for latency-sensitive applications that require data delivery to be timely, efficient and reliable. This trend is primarily driven by the proliferation of wireless networks of real-time data-gathering sensor-actuator devices. This has led to a strong need to bring real-time concerns to the forefront of an integrated research thrust into wireless real-time systems. In this paper, we introduce and analyze a specific instance of the rich set of problems in this domain. We consider a wireless network serving real-time flows in which the underlying physical layer provides multiple transmission rates. Higher rates have more stringent SINR requirements and thus represent a trade-off between raw transmission speed and packet error rate. We adopt a first principles approach to the design of optimal real-time scheduling algorithms for such a multi-rate wireless network. We illustrate the inherent complexities of the problem through examples and obtain provably optimal structural results. We then characterize the optimal policy for an approximate model. Our theoretical analysis provides guidelines for heuristic scheduler design. Our initial work indicates that this is a rich problem domain with the potential for a unifying theory that integrates real-time requirements into multi-rate wireless network design.

## Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Network Architecture and Design—*Wireless communication*; C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems

## General Terms

Performance, Theory

## Keywords

Real-time scheduling, wireless networks, rate adaptation, dynamic programming.

## 1. INTRODUCTION

The rapid proliferation of wireless sensor-actuator networks has brought to a fore the need to address real-time concerns in wireless networking. Such real-time applications do not merely require efficient and reliable data delivery. Timeliness is of utmost importance, as data communication is often part of a complex sequence of time-critical actions performed by the sensor network in response to external events. Real-time guarantees may also be required for multimedia applications in wireless mesh networks.

Typical wireless systems provide multiple transmission rates, e.g., 1, 2, 5.5 and 11 Mbps in IEEE 802.11b, 6, 9, 12, 18, 24, 36, 48, 54 Mbps in IEEE 802.11g and 11, 22, 33, 44, 55 Mbps in IEEE 802.15.3. Higher rates have more stringent SINR requirements for decoding, and represent a trade-off between raw transmission speed and packet error rate. In this work, we analyze the real-time implications of this trade-off using a first principles approach. We study a canonical problem, focusing on the case where all flows are over a single wireless link. Given the probabilistic nature of wireless losses, we take a *soft-real-time* approach and focus on minimizing the expected deadline-miss-ratio. We use illustrative examples to demonstrate the complex set of issues introduced by this seemingly simple problem. We show that intuitive solutions, e.g., earliest-deadline-first (EDF) scheduling at the rate with the minimum expected transmission time (min-ETT), are

---

[*] Supported by a Vodafone Graduate Fellowship.

[†] Supported by a Vodafone Graduate Fellowship.

[‡] This research was supported in part by the National Science Foundation through grant CNS-0613665.

*MobiCom'07*, September 9–14, 2007, Montréal, Québec, Canada.
ACM 978-1-59593-681-3/07/0009.

not necessarily optimal. We also show that packet scheduling and rate selection are inextricably linked, and even apparently simple scenarios are computationally intractable.

Motivated by the intractability of the problem, we adopt a two pronged approach toward developing theoretical insight. Firstly, **we establish the structure of the optimal policy** by showing that it is possible to associate each packet with a "pseudo-deadline" (a value between its arrival-time and deadline) such that the optimal scheduler has an earliest pseudo-deadline first (EPDF) structure. Secondly, **we consider an approximate, albeit more tractable model**. For this model, we show that the "packet arbitration problem", i.e., which packet to schedule, is decoupled from the "rate selection problem", i.e.., what rate to schedule the packet at. We further establish that EDF is the optimal packet arbitration policy, and min-ETT is the optimal rate selection rule.

Our work indicates the potential for a unifying theory that integrates real-time requirements into multi-rate wireless network design.

## 2. MODEL/NOTATION/TERMINOLOGY

We study a slotted-time, single-link wireless system, where all flows have the same source-destination pair. We model a multi-rate wireless system, with each rate corresponding to a different modulation, e.g., BPSK, QPSK, 16-QAM and 64-QAM in IEEE 802.11b. Faster rates have higher SINR requirements for successful decoding, and thus exhibit higher packet error rates. There are $k$ rates, labeled as $r_1, ..., r_k$, where each $r_i$ is specified by a two-tuple $(l_i, p_i)$. A packet transmission at rate $r_i$ takes $l_i$ time slots, and is successful with probability $1 - p_i$. All $l_i$'s are assumed integral. Thus, we implicitly assume that all packets have the same size.

We denote absolute deadlines by $d$, and relative deadlines by $D$. We consider two packet arrival models:

1. *One shot model*: $n$ flows, comprising one packet each, arrive at $t = 0$ with deadlines $d_1 < d_2 < ... < d_N$. No arrivals occur after $t = 0$.
2. *Periodic model*: For flow $i$, packets arrive periodically with period $T_i$ with initial relative-deadline $D_i = T_i$. The hyper-period is the least common multiple(LCM) of all $T_i$'s.

Traditional real-time schedulability analysis focuses on finding conditions on arrival rates so that all flows are guaranteed to meet their deadlines. In a system with probabilistic losses, this is an impossible undertaking, and we focus on minimizing the deadline miss ratio.

## 3. A FIRST GLANCE AT THE PROBLEM

We begin with some observations demonstrating that many apparently intuitive properties do not hold for the optimal scheduling policy.

EXAMPLE 1. *Consider a single one-shot flow with $d = 4$ over a link with two rates $r_1 = (2, p_1)$ and $r_2 = (3, p_2)$, with $p_1^2 > p_2$, e.g., $p_1 = 0.5$, $p_2 = 0.2$. Then, the optimal rate sequence is $r_2$, which yields a deadline miss probability of $p_2$, but leaves one slot unused, even though the packet's deadline is yet to expire.*

This yields:

OBSERVATION 1. *Unless there is a rate $r_i$ with $l_i = 1$, the optimal scheduler may not be work-conserving.*

Similarly, EDF scheduling may seem intuitively optimal. This is not always true, as shown by the following counter-example:

EXAMPLE 2. One-shot arrivals: *Consider two flows $\tau_1$ and $\tau_2$ with deadlines $d_1 = 1$ and $d_2 = 2$ respectively. There are two rates $r_1 = (1, 0.6), r_2 = (2, 0.1)$. Using EDF, we are forced to use $r_1$ in both slots, and incur $0.6 + 0.6 = 1.2$ expected deadline misses. If instead, we play only $\tau_2$ at rate $r_2$, the expected deadline misses would be $1 + 0.1 = 1.1$.*
Periodic arrivals: *Consider two flows $\tau_1$ and $\tau_2$ with periods $T_1 = 2$ and $T_2 = 4$ respectively. The hyper-period is thus of duration 4. There are two rates $r_1 = (1, 0.99), r_2 = (4, 0.01)$. Using any EDF ordering, the expected deadline misses per hyper-period is $2.96000001$. Instead, by playing only $\tau_2$ at rate $r_2$, we can reduce the expected deadline misses to $2 + 0.01 = 2.01$.*

Thus we obtain:

OBSERVATION 2. *EDF is not necessarily optimal for the multi-rate real-time scheduling problem.*

## 4. REALISTIC MODEL: CONSTANT SIZE PACKETS

We now build a theoretical understanding of real-time multi-rate scheduling for fixed size packets. Due to lack of space, results are stated without proof. For details, please see [2].

### 4.1 Single-Packet Rate Adaptation

Consider the apparently simple RATE-ADAPT problem, where a single packet arriving at $t = 0$ with deadline $d$, must be transmitted over a lossy multi-rate link. The link has a set of $k$ possible transmission rates $r_i = (l_i, p_i)$. We seek to find the sequence of transmission rates to be used, to minimize the probability of missing the deadline. An optimization instance of RATE-ADAPT can be represented as a 3-tuple $(d, \mathcal{P}_e, \mathcal{L})$, with $|\mathcal{P}_e| = |\mathcal{L}| = k$, $\mathcal{P}_e = \{p_1, p_2, ..., p_k\}$, and $\mathcal{L} = \{l_1, l_2, ..., l_k\}$.

THEOREM 1. *RATE-ADAPT is NP-hard.*

*Greedy Heuristic for RATE-ADAPT.* Consider the following greedy heuristic: sort the rates in non-decreasing order of $m_i = p_i^{1/l_i}$, and let the order be $r_{i_1}, r_{i_2}, ..., r_{i_k}$. We greedily include as many transmission attempts at $r_{i_1}$ as possible, then move on to $r_{i_2}$, and so on.

LEMMA 1. *If the packet transmission durations $l_i$ are* mutually harmonic*, i.e., for all $1 \leq i, j \leq k$, one of $l_i, l_j$ is a multiple of the other, then the greedy heuristic is optimal.*

LEMMA 2. *Let $p_{greedy}$ be the probability of deadline miss when using the heuristic. Let $p_{opt}$ be the optimal deadline miss probability. Then $p_{greedy} \leq \sqrt{p_{opt}}$.*

### 4.2 Multi-Rate Scheduling Problem

The multi-rate real-time scheduling problem MRATE-SCHED models the single-link scheduling situation described in Section 2. Given a flow-set $\tau$, we seek to determine the scheduling algorithm that minimizes the expected number of deadline misses. Note that **the problem is a joint packet and rate selection problem**.

LEMMA 3. *MRATE-SCHED is NP-hard.*

### *Structure of Optimal MRATE-SCHED Algorithm*

Theorem 2. *There exists an optimal MRATE-SCHED scheduling algorithm that has earliest pseudo-deadline first (EPDF) structure. If at some time $t$, this optimal scheduler bypasses the packet with the earliest deadline (call it $i$), and schedules another packet $j$, then $i$ will never be scheduled again by the optimal scheduler. This property holds for both one-shot and periodic arrival models. Additionally, the difference between a packet's deadline and pseudo-deadline is less than the maximum packet transmission time, over all possible rates.*

This structural result can be intuitively explained by observing that, since all packets experience the same channel conditions in the single-link problem, one can benefit from bypassing the packet with earliest deadline only if the packet we schedule instead is played at a slower and better rate that was not usable for the earliest deadline packet due to deadline constraints. When different packets see different channel conditions (as in a multi-link scenario), this intuition no longer holds, and it can be shown via counter-example that EPDF is not necessarily optimal.

## 4.3 Rate-Selection

In light of Theorem 2, one might wonder whether within a packet's *pseudo-deadline*, the optimal scheduler always selects the rate that is optimal for that individual packet. We show via counter-example that this is not the case.

Example 3. *Consider $2k$ periodic flows, each having period $T_i = T = 2k$ (this example may also be interpreted as an instance of the one-shot model). There are two available rates: $r_1 = (1, 0.4), r_2 = (k+1, 0.0)$. The hyper-period is $T$, and it suffices to consider this duration. From Theorem 2, there exists an optimal EPDF scheduler that bypasses a packet only if another packet can be scheduled at a slower (but better) rate. This implies that the pseudo-deadlines are the same as the deadlines for this example. Hence at the beginning of slot 1, one can schedule the first packet at the greedy best rate $r_2$ without exceeding its pseudo-deadline, leading to expected deadline misses of at least $k$. Instead, consider the following naive policy: try sending packet $i$ exactly once at rate $r_1$. This yields an expected loss of* $0.4(2k) = 0.8k < k$.

This yields the following observation:

Observation 3. *Rate-selection within a packet's pseudo-deadline cannot always be optimally performed independent of other unfinished packets in the system.*

Additionally, performance also depends on the *the order in which rates are attempted*, as shown by the following example:

Example 4. *There are two one-shot flows $\tau_1$ and $\tau_2$, with deadlines $d_1 = 3$ and $d_2 = 5$ respectively. There are three rates $r_1 = (1, p_1), r_2 = (2, p_2), r_3 = (4, p_3)$. We assume that the greedy ordering of rates is $r_3, r_2, r_1$, i.e., $p_1 \geq p_2^{\frac{1}{2}} \geq p_3^{\frac{1}{4}}$. Note that $\tau_1$ cannot use rate $r_3$. The possible rate-set in this case is $r_2, r_1$, and the greedy order of rates for $\tau_1$ is $r_2, r_1$. Suppose we use EDF, and $\tau_1$ is played in this order upto its deadline, followed by $\tau_2$. The expected number of deadline misses is $X = p_2 p_1 + (p_2 \cdot p_2 + (1-p_2) p_2 p_1)$. Now consider the alternate order $r_1, r_2$ for scheduling $\tau_1$. The expected number of deadline misses is now $Y = p_1 p_2 + (p_1 \cdot p_2 + (1-p_1) p_3))$. $X - Y = (p_2^2 - p_3)(1-p_1) \geq 0$ by our assumed greedy ordering. To illustrate, consider the following values of the $p_i$'s: $p_1 = 0.75, p_2 = 0.4, p_3 = 0.1$. These respect the greedy ordering, and provide a valid instantiation of this example. In this case, $X = 0.3 + 0.16 + 0.18 = 0.64$, $Y = 0.3 + 0.3 + 0.025 = 0.625$, and $X - Y = (p_2^2 - p_3)(1-p_1) = (0.16 - 0.1)(0.25) > 0$. Thus it is better to play the non-greedy rate order $r_1, r_2$ than the greedy sequence $r_2, r_1$.*

*One might ask whether the order-sensitivity is only an artefact of the potentially sub-optimal use of EDF in the above example. This is easily shown not to be the case. From Theorem 2, $\tau_1$'s pseudo-deadline can only be either of 1, 2 or 3. The last case is already covered in the EDF cases described above. Consider the first two cases. When pseudo-deadline is 1, $\tau_1$ can only be played once. Then the best expected loss is $Z = p_1 + p_3 = 0.75 + 0.1 = 0.85$. This is greater than $X, Y$ above. When the pseudo-deadline is 2, using rate $r_2$ for $\tau_1$ yields expected loss $W = p_2 + p_2 p_1 = 0.7$, and using $r_1, r_1$ yields expected loss $p_1^2 + p_1(p_2 p_1) + (1-p_1) p_3 = 0.5625 + 0.225 + 0.025 = 0.8125$. Thus the optimal scheduler has loss $Y$ and it uses a non-greedy ordering.*

This yields the following observation:

Observation 4. *The performance of the scheduling policy is sensitive to the order in which rates are attempted within the packet's pseudo-deadline.*

# 5. APPROXIMATION MODEL

Motivated by the intractability results, we have considered an approximate approach. We provide a brief summary of results here. For details, please see [2]. Consider the following channel model: suppose at time $t$, a flow is scheduled for transmission at rate $r_i$, and consider a small interval $(t, t + l_i \Delta t)$. Then, with probability $(1 - p_i)\Delta t$, the packet successfully completes in this time interval. This model ensures that when rate $r_i$ is used, the time between successful transmissions is exponentially distributed with parameter $l_i(1 - p_i)$, and is the continuous analogue of the geometric distribution between successful slots for the constant packet model, with the means of both distributions being equal.

We can formulate the problem of finding the optimal policy as a stochastic control problem using the theory of dynamic programming. For the one shot model, this is a finite time horizon problem. It is clear that the optimal policy never schedules a flow whose relative deadline is 0. Suppose then, that the relative deadlines are denoted by the vector $\bar{D} := (D_1, ..., D_N)$, with $0 < D_1 < ... < D_N$. A policy $\pi_t(\bar{D})$ is characterized by a two tuple $(flow_t(\bar{D}), r_t(\bar{D}))$ describing which flow and rate to schedule at time $t$. Let $V^{\pi}(\bar{D})$ be the total number of deadline misses incurred by $\pi$ starting with initial relative deadlines $\bar{D}$. We wish to find the stationary $\pi^*$ that minimizes $V^{\pi}(\bar{D})$. Observe that the optimal $V^*(\bar{D})$ satisfies the optimal substructure property of dynamic programming, i.e., any sub-schedule of the optimal schedule is also optimal.

For the periodic model, we focus on the expected deadline miss ratio per slot over the infinite time horizon $(0, \infty)$. Assuming for simplicity that there are two flows, we denote the relative deadlines as $D_1, D_2$. Define $x$ and $y$ as $0-1$ variables indicating whether the packet corresponding to the first and second flows has left the system (1) or not (0). Then, a policy $\pi_t(D_1, D_2, x, y)$ is characterized by the two tuple $(flow_t(D_1, D_2, x, y), r_t(D_1, D_2, x, y))$ describing what flow and transmission rate to use at time $t$. Define $V^{\pi}(D_1, D_2, x, y)$ as the expected average number of deadline misses starting from $(D_1, D_2, x, y)$. We wish to find the stationary $\pi^*(D_1, D_2, x, y)$ that minimizes the average number of deadline misses. Again, the optimal $V^*(D_1, D_2, x, y)$ satisfies an optimal substructure property, i.e., every subschedule of the optimal schedule is also optimal. [1]

## 5.1 One shot arrivals

THEOREM 3. *For the one-shot model, the optimal policy* $\pi^*(\bar{D})$ *that minimizes the total number of deadline misses has the following structure:*

1. *It schedules packets in earliest-deadline-first (EDF) order, i.e,* $flow^*(\bar{D}) = argmin_i D_i$.
2. *The optimal rate selection rule is independent of packet scheduling rule, i.e.,* $r^*(\bar{D}) = r^*$.
3. *It selects the rate with the minimum expected transmission time, i.e.,* $r^* = argmin_k \frac{l_k}{1-p_k}$.

## 5.2 Periodic arrivals

We use a discounted cost infinite horizon total cost formulation of the periodic model problem, where a deadline miss incurred $t$ units in the future is discounted by $e^{-\gamma t}$ with respect to a deadline miss incurred in the present. This ensures that the infinite horizon total cost is finite. An optimality result for the infinite horizon discounted cost formulation can be converted into an optimality result for the average cost formulation, which is the deadline miss ratio formulation we are interested in, by taking the limit as $\gamma \to 0$ [1].

We have proved the following theorem for the periodic model with two flows [2]:

THEOREM 4. *The optimal policy* $\pi^*(D_1, D_2, x, y)$ *that minimizes the average number of deadline misses has the following structure:*

1. *When packets of both flows are in the system, it schedules packets in earliest-deadline-first order, i.e.,* $flow^*(D_1, D_2) = argmin_i D_i$. *When packets of only one of the flows is in the system, it schedules that flow's packet.*
2. *The optimal rate selection rule is independent of the packet scheduling rule, i.e.,* $r^*(D_1, D_2) = r^*$.
3. *The optimal transmission rate selection rule selects the rate with the minimum expected transmission time, i.e.,* $r^* = argmin_i \frac{l_i}{(1-p_i)}$.

OBSERVATION 5. *The multi-rate real-time scheduling problem in the approximate model has a closed form solution, and rate selection is decoupled from packet arbitration. This is in contrast to the realistic model with constant size packets, where the rate selection problem is, in general, NP-hard and cannot be decoupled from packet scheduling. This is true for both one shot and periodic arrivals.*

# 6. HEURISTIC APPROACHES

The theoretical insights derived in Sections 4 and 5 can be used to design heuristic policies. We have investigated the following rate-selection heuristics suggested by our theory:

1. *Select the rate that minimizes* $p_i^{1/l_i}$
2. *Select the rate with the minimum expected transmission time (ETT)* $\frac{l_i}{1-p_i}$
3. *Select the rate that minimizes a hybrid metric*: $\alpha p_i^{1/l_i} + (1-\alpha)\frac{ETT_i}{\max_i ETT_i}$

In terms of joint rate-and-packet selection policies, our theory suggests heuristics such as instantaneous greedy choice policies, and partial lookahead policies.

# 7. RELATED WORK

The seminal work of Liu and Layland [6] in real-time scheduling established the optimality of the earliest-deadline-first (EDF) scheduling policy.

Many rate adaptation algorithms have focused on boosting throughput performance. Of these, SampleRate [3] has most relevance to our work because it addresses the dichotomy between transmission rates and packet error rate, albeit from a throughput maximization perspective.

[7] describe a first principles approach to real-time scheduling in wireless environments with multi-user diversity, and use stochastic dynamic programming to show a "virtual-deadline-first" structure for the optimal policy. [4] have developed a similar dynamic programming approach for an inter-packet deadline (IPD) arrival model, and shown a switching curve structure for the optimal policy. A simulation-based approach is described in [5] for scheduling packets with deadlines in multi-rate wireless networks under a "continuum of rates" assumption.

# 8. REFERENCES

[1] D. Bertsekas. *Dynamic Programming and Optimal Control: Vol I and II.* Athena Scientific, 2005.

[2] V. Bhandari, V. Raghunathan, B. D. Bui, and M. Caccamo. Real-time implications of multiple transmission rates in wireless networks. Technical report, University of Illinois, Urbana-Champaign, 2007.

[3] J. Bicket. Bit-rate selection in wireless networks. Master's thesis, MIT, 2005.

[4] A. Dua and N. Bambos. Deadline constrained packet scheduling in wireless networks. In *Proc. of IEEE VTC 2005.*

[5] K. M. Elsayed and A. K. Khattab. Channel-aware earliest deadline due fair scheduling for wireless multimedia networks. *Wirel. Pers. Commun.*, 38(2):233–252, 2006.

[6] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM*, 20(1), 1973.

[7] V. Raghunathan, M. Cao, and P. R. Kumar. Real-time scheduling for wireless systems with losses. Technical report, University of Illinois, Urbana Champaign, 2007.

---

[1] There is a technical difference between the two formulations; the one shot model is a finite horizon problem, while the periodic model is an infinite horizon problem.

[2] We can generalize the periodic model proof to the case with $N$ periodic flows; the calculations are more involved in this case