

Weighted Multi-ID Allocation for Quality-of-Service Provisioning in 10BASE-T1S PLCA-Based In-Vehicle Networks

Minsu Kim, Yeonho Kang, Sehwa Jung, Hanyoung Park, and Ji-Woong Choi

Department of Electrical Engineering and Computer Science, DGIST, Daegu, South Korea

Email: {excel2001, dusgh7227, boxxerman, prkhnyng, jwchoi}@dgist.ac.kr

Abstract—10BASE-T1S is a standardized single-pair multidrop Ethernet technology for cost-effective intra-zone in-vehicle networks. However, conventional one-ID-per-node PLCA provides equal cyclic access and does not account for heterogeneous deadlines and traffic loads. This paper proposes a standards-compatible weighted multi-ID allocation method for quality-of-service (QoS) provisioning in PLCA-based networks. The method allocates logical PLCA IDs according to deadline urgency and traffic pressure, and places them using a largest-gap rule while preserving the standardized PLCA cyclic access procedure and Ethernet frame format. In a high-load mixed-criticality stress scenario, the proposed method uses 37.5% fewer IDs than a fixed-budget baseline, reduces the average critical-node deadline miss ratio from 19.03% to 5.64%, and avoids the deadline violation of the high-load node observed in the baseline.

Index Terms—10BASE-T1S, Physical Layer Collision Avoidance (PLCA), in-vehicle networks, Automotive Ethernet, Quality of Service (QoS) provisioning, multi-ID allocation, mixed-criticality traffic

I. INTRODUCTION

Modern in-vehicle networks (IVNs) are required to support heterogeneous traffic from various applications, such as advanced driver-assistance systems (ADAS), autonomous driving functions, chassis control, powertrain control, radar, and infotainment [1], [2]. These traffic flows differ in payload size, generation period, and latency requirements. In particular, safety-critical and control-related messages require low latency and predictable delivery, whereas sensor and multimedia traffic can impose high communication loads because of large payloads or short generation periods [2]. Meanwhile, automotive electrical/electronic (E/E) architectures are moving from domain-based architectures toward zonal architectures to reduce wiring complexity, vehicle weight, and system cost while improving scalability [3]. In a zonal architecture, many low- and medium-speed edge nodes must be connected within each zone, and conventional point-to-point Ethernet can increase the number of cables, switch ports, and transceivers. Therefore, 10BASE-T1S is attractive for intra-zone communication because it provides a standardized single-pair, half-duplex, multidrop Ethernet PHY, allowing multiple edge nodes to share a common bus while remaining compatible with the Ethernet ecosystem [4], [5].

The deterministic use of this shared bus is enabled by Physical Layer Collision Avoidance (PLCA), which provides

collision-free cyclic access by assigning transmission opportunities (TOs) according to a configured PLCA ID sequence [4]. However, the conventional one-ID-per-node PLCA configuration gives each node one transmission opportunity per cycle, regardless of its traffic deadline or offered load. This equal-access behavior is simple and predictable, but it is not well suited to mixed-criticality IVN traffic. A tight-deadline control node may need more frequent access opportunities to reduce waiting time, while a high-load node with large or frequent packets may also require sufficient service opportunities to avoid queue buildup. Therefore, quality-of-service (QoS) differentiation in 10BASE-T1S should consider not only deadline urgency but also the offered communication demand of each node.

Several approaches have been studied to improve latency or QoS in 10BASE-T1S networks, including node prioritization and packet segmentation [6], interrupt-enabled PLCA [7], and scheduling mechanisms for multidrop Ethernet environments [8], [9]. These methods can improve the delay performance of delay-sensitive traffic, but they generally require modifications or extensions to the basic PLCA operation, such as priority-aware access control, interrupt handling, or packet segmentation and reassembly. In contrast, PLCA ID assignment provides a configuration-level approach for differentiating service opportunities while preserving the standardized PLCA cyclic access procedure. Recent studies have shown that assigning multiple logical PLCA IDs to a single physical node can increase the node's service-opportunity frequency and enable QoS differentiation [10], [11].

However, existing multi-ID studies mainly demonstrate the feasibility of assigning multiple IDs and do not fully address how a limited ID budget should be allocated among nodes or how the assigned IDs should be placed within the PLCA cycle. A deadline-dominant allocation [10] can protect tight-deadline nodes, but it may concentrate too many IDs on a small number of critical nodes and leave high-load nodes underserved. Conversely, increasing the total number of IDs without considering placement can lengthen the PLCA cycle and degrade tail-delay behavior. These observations motivate a traffic-aware multi-ID configuration that jointly accounts for deadline urgency, traffic pressure, and ID placement.

To this end, we propose a weighted multi-ID allocation method for QoS differentiation in 10BASE-T1S PLCA-based

IVNs. The proposed method treats PLCA ID assignment as a standards-compatible offline configuration problem that determines both node-wise ID counts and their placement in the PLCA cycle while preserving the standardized PLCA cyclic access procedure and Ethernet frame format. It computes node weights from inverse-deadline urgency and offered-load traffic pressure, allocates PLCA IDs according to these weights, and uses a largest-gap insertion procedure to distribute the assigned IDs across the cycle. The main contributions of this paper are summarized as follows:

- We formulate PLCA multi-ID assignment as a standards-compatible offline QoS configuration problem that differentiates service opportunities through node-wise ID counts and PLCA ID sequencing.
- We propose a weighted multi-ID allocation method that jointly considers deadline urgency and traffic pressure, and we combine this allocation with a largest-gap insertion procedure to reduce the clustering of transmission opportunities within the PLCA cycle.
- We evaluate Basic PLCA, Burst mode, a fixed-budget Baseline Multi-ID scheme, and the proposed method using average waiting time, response time distribution, and packet-level deadline miss ratio. In a high-load mixed-criticality stress scenario, the proposed method reduces the average critical-node miss ratio from 19.03% to 5.64% while using 37.5% fewer PLCA IDs than the fixed-budget baseline.

II. SYSTEM MODEL AND PERFORMANCE METRICS

A. PLCA Multi-ID Model

We consider a 10BASE-T1S PLCA-based multidrop network in which N nodes share a common half-duplex physical medium. Each node maintains a local first-in-first-out (FIFO) transmit queue for packets generated by its traffic flow. In PLCA operation, a beacon initiates each cycle, after which TOs are granted sequentially according to the configured PLCA ID sequence. When the TO corresponding to an assigned ID arrives, the associated node transmits one frame if its local transmit queue is nonempty; otherwise, the TO is treated as an inactive opportunity. In the conventional one-ID-per-node configuration, each physical node obtains one TO per PLCA cycle. This provides simple round-robin access, but the TO frequency is identical for all nodes regardless of their traffic load or latency requirements. In contrast, a multi-ID configuration can associate multiple logical PLCA IDs with a single physical node. A node with more assigned IDs obtains more TOs within a PLCA cycle and can therefore access the shared medium more frequently.

Let m_i denote the number of PLCA IDs assigned to node i , and let K denote the total ID budget in one PLCA cycle. The ID allocation satisfies

$$m_i \geq 1, \quad \sum_{i=1}^N m_i = K. \quad (1)$$

The constraint $m_i \geq 1$ ensures that every node retains at least one service opportunity in each cycle, while larger m_i values

provide differentiated access frequency. The PLCA access order within a cycle is represented by an ID sequence

$$\pi = [\pi_1, \pi_2, \dots, \pi_K], \quad (2)$$

where $\pi_k = i$ means that the k -th TO is assigned to node i . Thus, a PLCA ID configuration is specified by both the node-wise ID counts $\{m_i\}$ and the ID sequence π . The ID counts determine how many TOs each node receives, while the sequence determines when those TOs appear within the cycle.

B. Traffic and PLCA Overhead Model

We assume that each node generates a single data flow. Let $\bar{T}_i$ denote the average packet generation period of node i . The average packet arrival rate is given by

$$\lambda_i = \frac{1}{\bar{T}_i}. \quad (3)$$

When node i transmits a packet, its active service time is defined as

$$s_i^{\text{act}} = \frac{\eta \cdot 8(P_i + H_{\text{act}})}{R_b}, \quad (4)$$

where P_i is the payload size, R_b is the link bitrate, η is the 4B/5B encoding factor, and H_{act} is the fixed active-slot overhead including the Ethernet preamble, start-of-frame delimiter (SFD), Ethernet header, frame check sequence (FCS), and inter-frame gap (IFG). The active service time represents the channel occupancy when a packet is actually transmitted.

PLCA-specific overheads are modeled separately in the discrete-event simulator. These include the beacon interval at the beginning of each PLCA cycle, the silence interval of an inactive TO, and the burst timer interval used in Burst mode. Therefore, s_i^{act} is used to quantify the active transmission demand of each node, whereas the simulated packet response time includes both active transmission time and PLCA access delays caused by other nodes, inactive opportunities, and cycle-level overheads. To capture the active-channel demand of each node, we use the offered-load form commonly used in queueing and network-capacity analysis. Specifically, the traffic pressure of node i is defined as

$$\rho_i = \lambda_i s_i^{\text{act}}. \quad (5)$$

Note that ρ_i increases when packets are generated more frequently or when the active transmission time becomes longer. Thus, it captures the offered communication demand imposed by both packet generation period and payload size. This is important because a node can impose high communication demand through a large payload or a short generation period, even if it is not the most delay-critical flow.

C. Delay and Deadline-Miss Metrics

Let D_i denote the relative deadline of each packet generated by node i , measured from its arrival time. For the p -th packet of node i , the waiting time is defined as

$$w_{i,p} = b_{i,p} - a_{i,p}, \quad (6)$$

where $a_{i,p}$ and $b_{i,p}$ denote the packet arrival time and transmission start time, respectively. The packet response time is given by

$$T_{i,p}^{\text{resp}} = w_{i,p} + s_i^{\text{act}}. \quad (7)$$

Note that average waiting time is used to compare the overall delay tendency of each configuration. However, deadline satisfaction must be evaluated at the packet level because a small average delay does not necessarily guarantee that all packets meet their deadlines. To compare packets with different deadlines, we define the deadline-normalized response ratio as

$$\delta_{i,p} = \frac{T_{i,p}^{\text{resp}}}{D_i}, \quad (8)$$

and it implies that a packet satisfies its deadline when $\delta_{i,p} \leq 1$, whereas $\delta_{i,p} > 1$ indicates a deadline miss. The deadline miss ratio of node i is defined as

$$\epsilon_i = \frac{1}{N_i^{\text{gen}}} \sum_{p=1}^{N_i^{\text{gen}}} \mathbf{I}(T_{i,p}^{\text{comp}} - a_{i,p} > D_i), \quad (9)$$

where N_i^{gen} is the number of packets generated by node i during the measurement interval, $T_{i,p}^{\text{comp}}$ denotes the completion time of packet p , and $\mathbf{I}(\cdot)$ is the indicator function. Packets remaining in the queue after their deadlines expire are also counted as deadline misses.

III. PROPOSED WEIGHTED MULTI-ID ALLOCATION

The proposed configuration procedure has two levels. For a fixed candidate ID budget K , Sections III-A and III-B generate one PLCA configuration: Section III-A allocates the node-wise ID counts, and Section III-B constructs the corresponding ID sequence. Section III-C then performs the outer design-time search over candidate ID budgets by simulating each generated configuration and selecting the final budget K^* .

A. Weight-Based ID Count Allocation

For a fixed candidate budget K , the proposed method first computes node priority weights and allocates the node-wise ID counts. Jointly optimizing the ID budget K , the ID counts $\{m_i\}$, and the ID sequence π under packet-level deadline constraints is a combinatorial problem. Therefore, the proposed method decomposes the configuration problem into priority-weight calculation, ID count allocation, and ID sequence construction.

Since deadline urgency and traffic pressure have different units and scales, we normalize both terms by their maximum values across nodes. The normalized deadline urgency and normalized traffic pressure of node i are defined as

$$u_i = \frac{1/D_i}{\max_j (1/D_j)}, \quad p_i = \frac{\rho_i}{\max_j \rho_j}. \quad (10)$$

The priority weight of node i is then defined as

$$\omega_i = \gamma_{\text{dl}} u_i^{\gamma_{\text{exp}}} + \gamma_{\text{tr}} p_i, \quad (11)$$

where γ_{dl} and γ_{tr} control the relative importance of deadline urgency and traffic pressure, respectively. We use a sublinear

exponent $\gamma_{\text{exp}} < 1$ for the deadline-urgency term to maintain high priority for tight-deadline nodes while reducing excessive ID concentration on a small set of critical nodes.

After computing the priority weights, all nodes are initialized with one PLCA ID, i.e., $m_i = 1$. The remaining $(K - N)$ IDs are assigned using a weighted virtual-counter rule, inspired by deterministic proportional-share scheduling [12]. The weights are first normalized as

$$\tilde{\omega}_i = \frac{\omega_i}{\sum_{j=1}^N \omega_j}. \quad (12)$$

Subsequently, the virtual counter of node i is initialized as $v_i = 1/\tilde{\omega}_i$. At each allocation step, one additional ID is assigned to the node with the smallest virtual counter:

$$i^* = \arg \min_i v_i \quad (13)$$

$$m_{i^*} \leftarrow m_{i^*} + 1, \quad v_{i^*} \leftarrow v_{i^*} + \frac{1}{\tilde{\omega}_{i^*}}. \quad (14)$$

Thus, a node with a larger priority weight has a smaller counter increment and is therefore selected more frequently.

B. Largest-Gap ID Sequence Construction

After obtaining the ID counts for a candidate budget K , the proposed method constructs the corresponding PLCA ID sequence $\pi(K)$ as a circular list of K slots. A slot is called empty if no PLCA ID has been assigned to it. A global gap denotes a maximal consecutive segment of empty slots in the circular sequence. For a node that already has assigned IDs, its current gap denotes the largest circular interval between two consecutive IDs of the same node.

Nodes are processed in descending order of ω_i , with ties resolved by node index. The first ID of each node is placed at the midpoint of the largest global gap. For the remaining IDs of the same node, the target position is chosen as the midpoint of the node's largest current gap. If the target position is already occupied, the nearest empty slot is used. This largest-gap rule spreads repeated IDs of the same node over the cycle and reduces long inter-opportunity gaps.

C. ID-Budget Selection

The proposed method selects the total ID budget K from a finite candidate set $\mathcal{S}_K = \{N, N+1, \dots, K_{\text{max}}\}$, where N is the number of physical nodes and K_{max} is the maximum ID budget considered in the design-time search. For each $K \in \mathcal{S}_K$, the PLCA configuration generated by Sections III-A and III-B is evaluated by design-time simulation to estimate the average response time $\bar{T}_i^{\text{resp}}(K)$ of each node. A candidate is feasible if

$$\bar{T}_i^{\text{resp}}(K) \leq D_i, \quad \forall i. \quad (15)$$

Among feasible candidates, the selected budget minimizes the average response time of the critical-node set $\mathcal{C}$:

$$K^* = \arg \min_{K \in \mathcal{S}_K} \frac{1}{|\mathcal{C}|} \sum_{i \in \mathcal{C}} \bar{T}_i^{\text{resp}}(K), \quad \text{s.t. } \bar{T}_i^{\text{resp}}(K) \leq D_i, \quad \forall i. \quad (16)$$

If no candidate is feasible, the candidate with the smallest maximum normalized average response time is selected.

TABLE I
TRAFFIC SPECIFICATION OF NODES

Node	Type	Payload [B]	Period [ms]	Deadline [ms]
1	Powertrain	46	5	5
2	Safety	46	1	1
3	Chassis	46	1	1
4	Multimedia	700	2	5
5	Passenger	46	100	100
6	Comfort	46	250	250
7	Diagnostics	100	100	100
8	Radar	750	2	2

IV. PERFORMANCE EVALUATION

A. Simulation Setup

To evaluate the proposed method, we implement a MATLAB-based discrete-event PLCA simulator. The simulator uses a simplified PLCA timing model that explicitly accounts for the beacon interval, active transmission time, inactive silence interval, and burst timer interval. We simulate a 10BASE-T1S multidrop network with $N = 8$ nodes sharing a single half-duplex link. The link bitrate is set to $R_b = 10$ Mbps, and the 4B/5B encoding overhead is reflected by using $\eta = 1.25$. A 20-bit beacon is inserted at the beginning of each PLCA cycle. Following the Ethernet frame format and inter-frame gap specified in IEEE 802.3, the fixed active-slot overhead is set to 38 bytes, consisting of the 7-byte preamble, 1-byte SFD, 14-byte Ethernet header, 4-byte FCS, and 12-byte IFG [4].

The traffic specification is summarized in Table I. Nodes 2 and 3 are configured as critical nodes with a 1 ms period and a 1 ms deadline, whereas nodes 4 and 8 represent high-load traffic with large payloads and short generation periods. The offered load is approximately 0.95 when active-slot overhead is included. We choose this near-saturation load to evaluate the robustness of PLCA ID allocation under high contention, following high-load bus-topology IVN evaluation scenarios such as [5]. Under lighter loads, most nodes can meet their deadlines even with simple round-robin access, which can mask the effect of service-opportunity differentiation.

To compare the effects of different PLCA service-opportunity configurations, we evaluate four modes. Basic PLCA assigns one PLCA ID to each node. Burst mode also uses one ID per node, but allows a node to transmit up to two packets consecutively within one TO. Baseline Multi-ID follows the deadline-driven multi-ID configuration reported in [10]. It fixes the total ID budget to $K = 32$ and assigns additional IDs according to a deadline-driven cycle-pressure metric. This baseline represents a fixed-budget, critical-node-oriented use of multiple PLCA IDs. Unlike the proposed method, it does not explicitly account for offered traffic pressure.

In contrast, the proposed method performs the design-time K -selection procedure described in Section III and selects $K = 20$. Therefore, the proposed method differs from the baseline in two aspects: it searches for a smaller feasible ID budget and allocates IDs using both deadline urgency and traffic pressure. For the proposed method, we set $\gamma_{dl} = \gamma_{tr} = 1$,




Fig. 1. Node-wise average waiting time under the four PLCA configurations.

$\gamma_{exp} = 0.5$, $\mathcal{C} = \{2, 3\}$, and $\mathcal{S}_K = \{8, \dots, 64\}$. The exponent $\gamma_{exp} = 0.5$ is used as a sublinear design choice to avoid over-concentration of IDs on the tightest-deadline nodes. The design-time K -search uses a 1 s simulation, and the selected configuration is evaluated over a 5 s simulation interval.

B. Results and Analysis

Fig. 1 compares the average waiting time of each node. Basic PLCA gives all nodes the same TO frequency and therefore causes severe queue accumulation at critical nodes 2 and 3. Burst mode reduces the average delay by allowing consecutive packet transmissions within one TO, but it does not increase the TO frequency of critical nodes. Baseline Multi-ID reduces the delay of nodes 2 and 3 by assigning many IDs to them under $K = 32$. However, this deadline-driven allocation provides insufficient service to node 4, which has a 700-byte payload and a 2 ms period. Although nodes 4 and 8 both generate high-load traffic, the deadline-driven baseline favors node 8 because of its tighter deadline, leaving node 4 underserved despite its high offered load. By incorporating traffic pressure, the proposed method mitigates this imbalance.

Fig. 2 shows the packet response-time distribution of each node. The baseline improves the response-time distribution of the critical nodes compared with Basic PLCA, but it causes a large response-time tail at node 4. The proposed method provides a more balanced response-time distribution across critical and high-load nodes, indicating that the improvement is not limited to average waiting time.

Table II summarizes the packet-level deadline miss ratios. Basic PLCA results in miss ratios of 99.71% and 99.87% for nodes 2 and 3, respectively. Burst mode reduces these values, but still leaves miss ratios of 25.14% and 41.25%. The Baseline Multi-ID scheme reduces the critical-node miss ratios to 17.94% and 20.12%, but this comes at the cost of a 95.39% miss ratio at node 4. By contrast, the proposed method selects $K = 20$, corresponding to 37.5% fewer IDs than the baseline, and reduces the miss ratios of nodes 2 and 3 to 4.69% and 6.58%, respectively, while keeping the miss ratio of node 4 at 0%. Overall, the proposed method reduces the average critical-node miss ratio from 19.03% to 5.64% compared with the baseline, corresponding to a relative reduction of approximately 70.4%. These results show that a larger fixed




Fig. 2. Node-wise packet response time distribution under the four PLCA configurations.

TABLE II

PACKET-LEVEL DEADLINE MISS RATIO OF CRITICAL AND HIGH-LOAD NODES

Mode	ID budget	Node 2	Node 3	Crit. Avg.	Node 4
Basic PLCA	8	99.71%	99.87%	99.79%	0.00%
Burst Mode	8	25.14%	41.25%	33.20%	0.00%
Baseline	32	17.94%	20.12%	19.03%	95.39%
Weighted Multi-ID	20	4.69%	6.58%	5.64%	0.00%

ID budget with deadline-only prioritization is not necessarily sufficient for mixed-criticality QoS. The improvement comes from the proposed deadline-and-traffic-aware allocation, which distributes the limited ID budget by considering both critical-node urgency and high-load service demand.

V. CONCLUSION

This paper proposed a standards-compatible weighted multi-ID allocation method for QoS differentiation in 10BASE-T1S PLCA-based in-vehicle networks. By jointly considering deadline urgency and traffic pressure, the proposed method configures node-wise PLCA ID counts and their placement while preserving the standardized PLCA cyclic access procedure and Ethernet frame format. Simulation results showed that the proposed method provides a more balanced PLCA configuration than a deadline-dominant fixed-budget baseline, reducing critical-node deadline misses and avoiding under-service of high-load nodes while using fewer PLCA IDs.

ACKNOWLEDGMENT

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) under the project No. RS-2024-00398157, "Development of 6G System Technologies Supporting AI-Native Application Services."

REFERENCES

- [1] V. K. Kukkala, J. Tunnell, S. Pasricha, and T. H. Bradley, "Advanced driver-assistance systems: A path toward autonomous vehicles," *IEEE Consumer Electronics Magazine*, vol. 7, no. 5, pp. 18–25, Sep. 2018.
- [2] H. Park, M. Cho, Y. Jang, and J.-W. Choi, "Latency analysis of in-vehicle network for advanced driver assistance system," in *Proc. IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, Washington, DC, USA, Oct. 2024, pp. 1–7.
- [3] H. Zhu, W. Zhou, Z. Li, L. Li, and T. Huang, "Requirements-Driven Automotive Electrical/Electronic Architecture: A Survey and Prospective Trends," *IEEE Access*, vol. 9, pp. 100 096–100 112, 2021.
- [4] *IEEE Standard for Ethernet—Amendment 5: Physical Layer Specifications and Management Parameters for 10 Mb/s Operation and Associated Power Delivery over a Single Balanced Pair of Conductors*, IEEE Std. 802.3cg-2019, 2020.
- [5] E. Choi, H. Song, S. Kang, and J.-W. Choi, "High-Speed, Low-Latency In-Vehicle Network Based on the Bus Topology for Autonomous Vehicles: Automotive Networking and Applications," *IEEE Vehicular Technology Magazine*, vol. 17, no. 1, pp. 74–84, Mar. 2022.
- [6] J. Min and Y. Park, "Performance Enhancement of In-Vehicle 10BASE-T1S Ethernet Using Node Prioritization and Packet Segmentation," *IEEE Access*, vol. 10, pp. 103 286–103 295, 2022.
- [7] —, "Interrupt-Enabled Physical Layer Collision Avoidance for 10BASE-T1S Automotive Ethernet," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 5, pp. 6083–6092, May 2024.
- [8] D. A. Nascimento, S. Bondorf, and D. R. Campelo, "Modeling and Analysis of Time-Aware Shaper on Half-Duplex Ethernet PLCA Multidrop," *IEEE Transactions on Communications*, vol. 71, no. 4, pp. 2216–2229, Apr. 2023.
- [9] M. Elattar, D. Minagias, Z. Lou, and G. Carle, "A Priority-Based Scheduling Mechanism for Multidrop Networks," in *Proc. IEEE Global Communications Conference (GLOBECOM)*, Taipei, Taiwan, Dec. 2025, pp. 4601–4606.
- [10] F. Cerruto, L. Leonardi, L. Lo Bello, and G. Patti, "Performance Assessment of Ethernet 10BASE-T1S Networks in Automotive Applications," in *Proc. IEEE 101st Vehicular Technology Conference (VTC2025-Spring)*, Jun. 2025, pp. 1–5.
- [11] —, "Achieving Quality of Service Differentiation in Ethernet 10BASE-T1S Nodes," in *Proc. IEEE 30th International Conference on Emerging Technologies and Factory Automation (ETFA)*, Porto, Portugal, Sep. 2025, pp. 1–4.
- [12] C. A. Waldspurger and W. E. Weihl, "Stride scheduling: Deterministic proportional-share resource management," MIT Laboratory for Computer Science, Tech. Rep. MIT/LCS/TM-528, 1995.