An Asymptotically Optimal Greedy Algorithm for Large Optical Burst Switching Systems

Lachlan L. H. Andrew* Yuliy Baryshnikov† E. G. Coffman, Jr.‡ Stephen V. Hanly
Jolyon White

ABSTRACT

As the number of wavelengths in OBS systems increases, the utilization achievable for a given blocking probability can be made to approach 100%. This paper shows that this property applies to a wavelength allocation algorithm of greedy type. Another property of this rule, one shared by most other wavelength assignment algorithms, is that, since lost traffic tends to occur near destinations, where the resource usage wasted by such traffic is large, very low blocking probabilities are important for efficient operation. To help identify regions of low blocking probability, we derive an asymptotically exact condition for zero blocking probabilities; it has a form reminiscent of the stability condition of the $M/G/1$ queue.

Keywords – Optical networks, Wavelength division multiplexing, Optical burst switching, Stochastic modeling, Fluid limits, Hydrodynamic limits.

1. INTRODUCTION

Optical burst switching (OBS) has attracted much attention recently [1, 2, 3, 4, 5, 6, 7, 8, 9] as a means of reducing the overhead associated with circuit setup on very high data rate all-optical networks. Although OBS is motivated by dense wavelength division multiplexing (DWDM), which is capable of carrying hundreds of wavelengths over a single fiber, few studies to date have considered systems with more than 32 wavelengths [8, 9]. This paper considers large systems, and shows that issues which are important for small systems, such as wavelength assignment, sort themselves out for realistic sized systems. The paper also discusses problems that arise from the disparity in service provided to bursts with short and long offsets.

OBS switches bursts of optical data which are long compared to the length of a packet, but may be short compared to the end-to-end delay and the electronic processing required for connection establishment. Before a burst is transmitted, a connection request is sent to alert switches on the route to establish a lightpath. Each switch is told the size of the burst and an *offset time* which specifies the advance notice, i.e., the delay until the arrival of the burst. When the burst is received at the switch, it passes straight through with no buffering and no conversion to electronic form. If the lightpath has not been established at the switch for any given burst arrival, the burst is simply dropped. Thus, blocked bursts can waste resources elsewhere along the lightpath where reservation requests were accepted. At the source, a burst is sent chasing after its (slower) connection request following a delay chosen so that the burst can not catch up with the request before the latter reaches its destination. The offset time therefore decreases on successive hops. As expected, this means that the failure to satisfy a connection request increases in probability as the request nears its destination. There are certain ways of dealing with this undesirable effect, e.g., see [11], but the problem emphasizes the need to keep blocking probabilities low.

When the connection request arrives, a switch must select the outgoing wavelength for the connection. The *Horizon* rule [10], which we analyze in this paper, gives arguably the simplest approach. As is typical, Horizon commits wavelengths at the times reservations are made. Such algorithms are more easily implemented than those not committing wavelengths until burst arrival times. A greedy rule of the latter type will be discussed briefly in the section containing final remarks.

Throughout the remainder of the paper, k denotes the number of wavelengths. Horizon partitions the current set of reserved intervals into subsequences, one for each wavelength, and is defined as follows.

Horizon algorithm: *Define the horizon $h_i(t)$, $1 \le i \le k$, at time t to be t if there are no currently reserved intervals on wavelength i; otherwise, it is the latest reserved time (right end of the rightmost reserved interval) on wavelength i. A request for interval (a,b) at time t is accepted and allocated to the wavelength with the largest horizon $h_i(t)$ no greater than a, if such a horizon exists. If no such horizon exists, the request is blocked. Allocation on wavelength i simply means adding the new interval as a new latest reserved interval to the state of wavelength i and updating $h_i(t)$ to b.*

A burst succeeds in reaching its destination only if all connection requests along the route have been accepted, but it makes the attempt irrespective of whether any of its requests have been blocked. Apart from this last property, each switch acts as a standard reservation system for connection requests; in the broader setting, the term "offset" is replaced by "advance notice." Note that Horizon does not attempt to exploit the voids it creates by using them for the intervals requested by later arrivals. Techniques for making use of voids include Just Enough Time (JET) [1] and Least Available Unused Channel with Void Filling (LUAC-VF) [5]. The obvious downside of such methods is the significantly greater time/space complexity of void-searching algorithms. Moreover, in the asymptotic limits studied here, these methods offer no improvement.

Indeed, a key observation of this paper is that the *relative* amount of void time per wavelength under Horizon tends to zero as the

*LLHA, SVH, and JW are all affiliated with the ARC Special Research Centre for Ultra-Broadband Information Neworks, Department of Electrical and Electronic Engineering, The University of Melbourne, Australia. Work funded in part by the Australian Research Council.

†Bell Labs, Lucent Technologies.

‡Electrical Engineering Department, Columbia University. Work funded in part by the National Science Foundation.

number of wavelengths increases while holding the per-wavelength traffic intensity fixed. In this limit, the capacity achievable by Horizon tends to that achievable by optimal scheduling. For several interesting network topologies, simulation results indicate that as few as 64 wavelengths may be enough for the asymptotic regime to give useful approximations to Horizon.

A mathematical model of the OBS system is presented next.

2. OBS MODEL

In our model of the Horizon algorithm at an OBS switch, the connection requests for transmission intervals arrive in a Poisson stream; the offset times τ and the durations d of the requests are independent of arrival times and are drawn from a given joint distribution. To avoid distractions, we assume that τ and d are independent, although much of our analysis would still apply were this assumption not to hold. For the same reason, we assume that the distribution of $\tau + d$ is strictly positive over a finite interval; the right endpoint of the interval is denoted by r^*. Finite support is not a restrictive assumption in applications, but the strictly-positive assumption entails approximations in a discrete world.

Our interest is in large-k asymptotics, more precisely the hydrodynamic limit, so the arrival rate scales with k. In particular, we take the arrival rate to be $k\lambda$ for given $\lambda > 0$, and thus keep the per-wavelength offered load constant as $k \to \infty$. A useful measure of the load at time t is given by the process

$$N^{(k)}(s,t) = \#\{i : h_i(t) > t + s\}$$

or its scaled, per-wavelength version,

$$n^{(k)}(s,t) = \frac{1}{k} N^{(k)}(s,t).$$

which we call the *occupancy* process. Thus, $n^{(k)}(s,t)$ is the fraction of the wavelengths to which a request arriving at time t with offset s cannot be assigned. As a concession to tractability, we study a continuous, differentiable hydrodynamic limit,

$$n^{(k)}(s,t) \to n(s,t)$$

as $k \to \infty$, and compute desired quantities in this deterministic limit as large-k approximations. The formalities concerning convergence to this limit are deferred to a more detailed report. Our objective here is to give convincing, albeit heuristic, arguments in support of the stated properties of the limit process.

The wavelengths at time $t + \Delta t$ with horizons exceeding $t + s$ consist of those with horizons exceeding $t + s + \Delta t$ at time t plus those assigned arrivals in $[t, t + \Delta t]$ with requested intervals that extend a horizon less than $t + s$ to a horizon greater than $t + s$. Let $f^{(k)}(s,t)$ be the probability that an arrival in $[t, t + \Delta t]$ is accepted and requests an interval extending a horizon from a point earlier than $t + s$ to a point later than $t + s$. Assuming that, in the hydrodynamic limit, $f^{(k)}(s,t)$ converges to a function $f(s,t)$ then we can write

$$n(s, t + \Delta t) = n(s + \Delta t, t) + \lambda \Delta t f(s,t) + o(\Delta t)$$

and therefore

$$\frac{\partial n(s,t)}{\partial t} = \frac{\partial n(s,t)}{\partial s} + \lambda f(s,t) \qquad (1)$$

For given boundary conditions, the problem comes down to finding the function $f(s,t)$. A workable approach is most clearly seen in the stationary regime, as in that regime we acquire a very useful monotonicity property. Eliminating the dependence on t in (1), we arrive at the simple ode

$$\frac{dn}{ds} = -\lambda f(s) \qquad (2)$$

where $n(\infty) = 0$ is enough to evaluate the constant of integration. We see immediately that the occupancy function $n(s)$ must be nonincreasing in s and that this is the limit of a standard stochastic ordering property of classical reservation systems: the greater the advance notice given, the greater the likelihood that a request is accepted. Precisely, at time t in the stochastic occupancy process, the fraction of reserved intervals covering $t + s$ is stochastically at most the fraction covering $t + s'$ for any $0 \le s' < s$. To see why this must hold, it is enough to observe that the arrival process is homogeneous in time, and that accepted requests with reserved intervals covering $t + s$ occurred in an interval $[t + s - r^*, t]$ which is properly contained in the interval $[t + s' - r^*, t]$ where requests with reserved intervals covering $t + s'$ arrived[1].

Since $n(s)$ is nonincreasing in s, if there is a "saturation" region of $\mathbf{R}^+$ where $n(s) = 1$, then it must be an initial interval $[0, s_*]$. Thus, $n(s)$ is constant at 1 throughout $[0, s_*]$, decreases monotonically[2] to 0 at $t + r^*$, and remains at 0 thereafter. Note that s_* can be regarded as the deterministic limit of the random thresholds $\min_{1 \le i \le k}\{h_i(t) - t\}$ in statistical equilibrium: arrivals with offsets less than the current threshold are blocked. Then blocking probabilities for requests with offsets less than s_* tend to 1 as $k \to \infty$.

Now consider what happens to void times as k becomes large. The horizon *profile* $(\hat{h}_i, i/k)$, $1 \le i \le k$, where the $\hat{h}_i$ are distributed as the $h_i(t) - t$ in decreasing size order, is a staircase function: when a new request is assigned a wavelength by Horizon, a void bounded by some step width is created. Step sizes decrease as k grows large; in the hydrodynamic limit, $k \to \infty$, the profile becomes a continuous trajectory $(\hat{h}(y), y)$, $0 \le y \le 1$, identical to the the occupancy trajectory $(s, n(s))$: If $s_* > 0$, the limiting profile has value 1 out to $\hat{h}(1) = s_*$ and from that point decreases monotonically to 0 at $\hat{h}(0) = r^*$ (i.e., $\hat{h}(y)$ increases monotonically from s_* to r^* as y decreases from 1 to 0). This continuous limit implies that the hydrodynamic limit of the equilibrium probability $f^{(k)}(s)$ that an arrival increases the count $N^{(k)}(s)$ is simply the joint probability that a requested interval is to the right of s_* and covers the point s, i.e.,

$$f(s) = \mathsf{P}(s_* < \tau < s < \tau + d)$$

Note that $f(s)$ depends on the arrival rate through the limit threshold s_*. When $s_* > 0$, continuity of the limit profile also implies that, under Horizon, the percentage of time wasted in voids tends to 0 as $k \to \infty$. The maximum rate at which reserved intervals expire on any wavelength is $1/\mathsf{E}[d]$, so

$$\lambda \mathsf{E}[d] > 1$$

is the condition for $s_* > 0$ in the hydrodynamic limit. In this case, the limiting arrival rate of accepted requests is $L(s_*)\lambda$, where $L(s) = \mathsf{P}(\tau > s)$ is the tail of the offset distribution. This arrival rate must equal the limiting departure rate $1/\mathsf{E}[d]$, and so s_* is the solution to $L(s) = (\lambda \mathsf{E}[d])^{-1}$ Thus, we have the following complete solution to (2).

If $\lambda \mathsf{E}[d] < 1$, then $s_* = 0$ and

$$n(s) = \lambda \int_s^{r^*} f(x)dx \qquad (3)$$

[1]The assumption of statistical equilibrium is needed here.

[2]Our assumption that the distribution of $\tau + d$ is strictly positive is at play here.

and if $\lambda \mathsf{E}[d] \geq 1$, then s_* is the solution to $L(s) = (\lambda \mathsf{E}[d])^{-1}$ and

$$n(s) = \begin{cases} 1, & 0 \leq s \leq s_* \\ \lambda \int_s^{r^*} f(x)dx & s > s_* \end{cases} \qquad (4)$$

Example. To illustrate the calculations, suppose τ and d have the uniform distribution on $[0,1]$. We compute f first, as follows. Assume $s_* > 0$. Then

$$f(x) = \begin{cases} \int_{s_*}^{x} (a+1-x)da = & (x-s_*)[1-\frac{1}{2}(x-s_*)], \\ & s_* < x \leq 1 \\ \int_{s_*}^{1} (a+1-x)da = & (1-s_*)[1-x+\frac{1}{2}(1+s_*)], \\ & 1 < x \leq s_*+1 \\ \int_{s_*+1}^{2} (a+1-x)da = & \frac{1}{2}(2-x)^2, \ s_*+1 < x \leq 2 \end{cases} \qquad (5)$$

As a partial check on these formulas, we find that

$$\int_{s_*}^{2} f(x)dx = 1/\lambda$$

gives, after breaking down the calculation as above, $(1-s_*)/2 = 1/\lambda$, which is the same as $L(s_*) = (\lambda \mathsf{E}[d])^{-1}$. In this example, we see that $s_* = 0$ if $\lambda < 2 = 1/\mathsf{E}[d]$. Substituting (5) into (4) gives $n(s)$, which is plotted in Figure 1 for the case $\lambda = 4$.

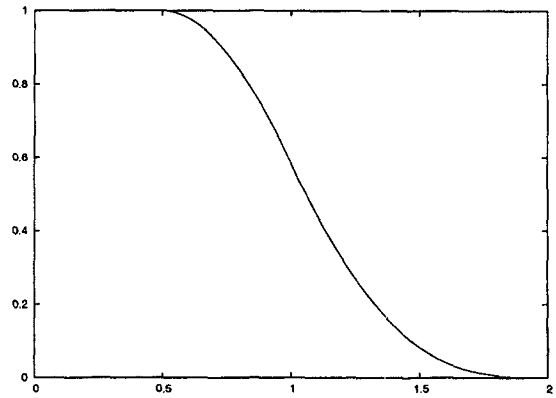

Figure 1: Limiting occupancy function n; $\lambda = 4$.

3. FINAL REMARKS

Another greedy algorithm interesting for its simplicity commits wavelengths only at burst arrival times. It is defined as follows. **Greedy algorithm:** *A connection request arriving at time t attempting to reserve an interval $[t+s, t+s+d]$ is accepted if and only if all points in the interval are covered by at most $k-1$ reserved intervals.*

If the request is accepted and the interval $[t+s, t+s+d]$ reserved, then at time $t+s$, a wavelength, say the j-th, will be made available to the corresponding burst throughout $[t+s, t+s+d]$. At time $t+s+d$, the reservation will be cleared from the system and the j-th wavelength returned to a pool of available wavelengths.

Greedy and Horizon have the same equilibrium hydrodynamic limit. However, it is not hard to see that Greedy is, implicitly, a void filling algorithm, so its large-k behavior requires a different argument than that for Horizon. We are currently writing up an analysis of the Greedy rule. We are also conducting an experimental comparison of Horizon and Greedy. To complete the analysis of both rules we are also deriving estimates of the total void time. Finally, we are investigating transient solutions $n(s,t)$, which appear to be significantly more elaborate than the results for the stationary case.

4. REFERENCES

[1] C. Qiao and M. Yoo, "Optical burst switching (OBS) – a new paradigm for an optical internet," *J. High Speed Networks*, vol. 8, pp. 69–84, Jan. 1999.

[2] L. Xu, H. Perros, and G. Rouskas, "Techniques for Optical Packet Switching and Optical Burst Switching," *IEEE Commun. Mag.*, pp. 136–142, Jan. 2001.

[3] M. Yoo, C. Qiao, and S. Dixit, "Optical burst switching for service differentiation in the next-generation optical internet," *IEEE Commun. Mag.*, vol. 39, pp. 98–104, Feb. 2001.

[4] M. Yoo, C. Qiao, and S. Dixit, "QoS performance of optical burst switching in IP-Over-WDM networks," *IEEE J. Selected Areas in Communications*, vol. 18, pp. 2062–2071, Oct. 2000.

[5] Y. Xiong, M. Vandenhoute, and H. Cankaya, "Control architecture in optical burst switched WDM networks," *IEEE J. Selected Areas in Communications*, vol. 18, pp. 1838–1851, Oct. 2000.

[6] M. Düser and P. Bayvel, "Performance of a dynamically wavelength-routed optical burst switched network," *IEEE. Photonic Technology Lett.*, vol. 14, pp. 239–241, Feb. 2002.

[7] A. Detti, V. Eramo, and M. Listanti, "Performance evaluation of a new technique for IP support in a WDM optical network: Optical composite burst switching (OCBS)," *IEEE J. Lightwave Tech.*, vol. 20, pp. 154–165, Feb 2002.

[8] K. Dolzer, C. Gauger, J. Spath, and S. Bodamer, "Evaluation of reservation mechanisms for optical burst switching," *Int. J. Electron. Commun.*, vol. 55, pp. 18–26, 2001.

[9] Z. Rosberg, H. Vu, and M. Zukerman, "Burst segmentation benefit in optical switching," *IEEE Comm. Lett.*, vol. 7, pp. 127–129, March 2003.

[10] J. Turner, "Terabit burst switching," *J. High Speed Networks*, vol. 8, pp. 3–16, Jan. 1999.

[11] B-C. Kim, J-H. Lee, Y-Z. Cho, and D. Montgomery, "A novel hop-by-hop priority increasing scheme for multi-hop optical burst switching networks," *Proc. Optical Internet and Photonic Switching Conference (COIN)*, pp. 87–89, July 2002, Cheju, Korea.