

Executive Summary

Abstract

The work performed under the aforementioned AFOSR grant includes research on satellite based communication networks. It consists of two parts. In the first part we investigate some of the arising resource allocation and handover issues in non-geostationary mobile satellite systems. We propose a channel classification scheme, in which available carriers are partitioned into classes and each class is associated with a range of propagation delays to the satellite. The key idea of our approach is that users with similar propagation delays are assigned to the same class of carriers. The suggested infrastructure results in better resource utilization and reduced call blocking rate and can be implemented with low signaling load. In the second part we consider information dissemination through satellite broadcast delivery and more specifically the identification and treatment of the issues that arise in novel multi-stage broadcast scheduling-caching applications. We investigate the design of the satellite schedule, the cache management policies at the ground stations and the terrestrial schedules that arise in joint fashion in the context of our problem. Our results demonstrate that such a two stage system is feasible and can provide more efficient data delivery compared to a single stage system.

1 Resource allocation and handover issues in non-geostationary mobile satellite networks

1.1 Resource allocation in MEO mobile satellite networks

In the near future, existing terrestrial radio networks are envisioned to integrate with satellite systems so as to provide global coverage. In order to establish communication for non-hand-held and hand-held user terminals, the radio link design must allow full- and half-duplex operation respectively, where the latter is desirable whenever radiation power restrictions are imposed. In addition, due to user mobility and volatility of the wireless medium, sophisticated resource management is required, so as to enhance system capacity and reduce the likelihood of blocking. An inherent problem of the satellite link is propagation delay, which can lead to inefficient resource allocation. This problem becomes more important in medium-earth-orbit (MEO) satellite systems, which are characterized by large propagation delays and large intra-beam delay variations.

The main contribution of this research initiative is to investigate some of the arising resource allocation and handover issues in non-geostationary mobile satellite systems. We propose a channel classification scheme, in which available carriers are partitioned into classes and each class is associated with a range of propagation delays to the satellite. The key idea of our approach is that users with similar propagation delays are assigned to the same class of carriers. The sug-

gested infrastructure results in better resource utilization and reduced call blocking rate and can be implemented with low signaling load.

1.2 System model

A satellite constellation in MEO orbit is considered. The projection of satellite on the earth is the sub-satellite point. The M are classified in L groups $\mathcal{B}_1, \dots, \mathcal{B}_L$. A group \mathcal{B}_i contains all equidistant beams from the sub-satellite point. Beams belonging to \mathcal{B}_i are type- i beams, for $i = 1, \dots, L$. Satellite gateways (GWs) constitute the access points for users. We use the terms "call", "user" and "user terminal (UT)" to refer to mobile users. An important parameter that characterizes the UT-satellite link is the satellite elevation angle θ with respect to the UT, which is the angle between the line that defines UT horizon and the line that connects the satellite to the UT.

A FDMA/TDMA access scheme is assumed. The available spectrum is divided into carrier frequencies. A TDMA frame in a carrier has 6 time slots, each of duration T_s , and a user traffic burst occupies one slot. A channel is perceived as a distinct carrier-slot pair. A UT transmits and receives once in a frame period. We assume half-duplex communication with non-overlapping transmission and reception intervals of a burst. A user can optionally use diversity within the same carrier. If a UT transmits and receives once in a frame and there is a guard time $t_g \ll T_s$ between transmission and reception intervals, 3 slots are required for single-channel half-duplex operation and 6 slots are needed in case of diversity. A static resource allocation scheme within each satellite is adopted and a set of carrier frequencies is assigned to each beam.

1.3 The delay class concept

Transmission (Tx) and reception (Rx) traffic bursts of the user are separated by a time offset, which is referred to as *Tx/Rx burst offset*. The relative positions of transmitted and received bursts in a window of 3 slots depend on Tx/Rx burst offset, which in turn depends on the UT propagation delay T_p to the satellite. Users in different positions in a beam require different Tx/Rx burst offsets to maintain non-overlapping transmission and reception intervals. If these users are assigned to the same carrier, several slots will remain unoccupied. In figure 1 the situation with a single or multiple offset values in a carrier are illustrated. For a single Tx/Rx burst offset, efficient call accommodation is achieved by packing users in contiguous time slots. However, for users with different delays, the system resorts to multiple burst offsets in order to maintain non-overlapping transmission and reception intervals and thus a significant amount of resources is wasted (in the figure, slots marked with "X"). For a given arrival rate, blocking rate is increased, since fewer calls are accommodated in the system.

In MEO mobile satellite systems, the situation above arises in edge beams with large intra-beam delay variations due to the curvature of earth surface. In order to circumvent the problem,

we introduce the concept of *delay class*. For a nominal delay T_0 , we define a delay range $\delta T_p \geq 0$ around T_0 . Transmission and reception bursts of UTs with propagation delays in the range $[T_0 - \delta T_p, T_0 + \delta T_p]$ are accommodated within a reference window. UTs with propagation delays in that range belong to the same *delay class* corresponding to T_0 and should be allocated to the same carrier or group of carriers. Several delay classes with delays $T_{0,i}$ must be defined and class i corresponds to a carrier group \mathcal{G}_i . Delays of users assigned to carriers of \mathcal{G}_i must satisfy $T_{0,i} - \delta T_p \leq T_p \leq T_{0,i} + \delta T_p$. Pictorially, each nominal delay $T_{0,i}$ corresponds to a contour (circle) on the earth surface. All nominal delays represent concentric circles, centered at the sub-satellite point Q (figure 2). A contour of delay $T_{0,i}$ comprises points on earth with that delay to the satellite. The contours of delay $T_{0,i} \pm \delta T_p$ form a zone that is defined as *delay class i*. Bursts of UTs in a delay class must arrive aligned at the satellite and UT [1]. Offset values for a specific UT are derived by comparing UT propagation delay T_p and nominal delay $T_{0,i}$ of its delay class and the offset is proportional to $T_p - T_{0,i}$. The range δT_p is derived by considering upper and lower bound of delays, such that accommodation in a window is feasible. It can be shown [2] that $\delta T_p = T_s/4 - t_g/2$.

Delay classes have certain positions with respect to the beam pattern. Figure 3 shows a projection of one quadrant of the beam pattern on a two-dimensional plane and the relative position of delay classes. Beams that are closer to the footprint edge are elliptical and elongated and therefore they cover a wide range of delays. A delay class serves a certain set of beams, namely beams of the same beam type \mathcal{B}_j due to circular symmetry. Beams of a beam type may be covered by more than one delay class. For example, out-most beams are covered by three delay classes, since propagation delay range is larger. Intermediate beams may be served by one or two delay classes.

Resources consist of carriers that belong to a pool and are assigned on a per beam basis with dynamic or fixed channel allocation (DCA, FCA) schemes. Although DCA schemes capture satellite movement and traffic variations, we consider fixed allocation of carriers to a beam for illustrative reasons. If arising calls are uniformly distributed in the area, the expected amount of traffic in out-most elongated beams is larger and call blocking is increased. The delay class concept can be used to alleviate this problem, by assigning more delay classes and therefore more carriers to these beams. The number of delay classes κ and their positions depend on satellite orbit and height, footprint and beam size, frame structure, traffic burst length, and guard time. Since adjacent delay classes may overlap, the maximum residence time at a delay class overlap area is another important design parameter that affects delay class positions. This residence time represents maximum allowed tolerance for delay class transitions.

Let ϕ be the earth central angle (figure 4). Let d and τ be the distance and delay from a delay class contour to satellite. Assume that H is the satellite height and R_E is the earth radius. A "cup" on the earth is the area defined by a delay class contour and a pole, and a "zone" is defined by delay class contours, i and $i + 1$. Let A be the area of such a cup or zone. The minimum and maximum elevation angles θ_{min} and θ_{max} stem from visibility conditions and definition of the closest delay class to satellite nadir. Range $[\theta_{min}, \theta_{max}]$ represents the portion of footprint to be covered with delay classes. Equivalently, this area is defined by central angle range $[\phi_{max}, \phi_{min}]$,

where the relation between angles ϕ and θ is

$$\phi(\theta) = \frac{\pi}{2} - \sin^{-1} \left(\frac{R_E}{R_E + H} \cos \theta \right) - \theta, \quad (1)$$

and $\phi_{min} = \phi(\theta_{max})$, $\phi_{max} = \phi(\theta_{min})$. The following algorithm finds delay class positions without considering delay class overlap regions.

- STEP 1 : Obtain the pair $(\phi_{min}, \theta_{max})$.
- STEP 2 : At first iteration, determine the first delay class. Compute parameter $h_0 = R_E (1 - \cos \theta_{max})$ (figure 4), i.e., the height of first (close to nadir) delay class. A cup is thus defined on earth surface at height h_0 and find the area of this cup, $A = 2\pi R_E h_0$. Then, compute distance parameters,

$$a_0^2 = R_E^2 - (R_E - h_0)^2 \text{ and } d_0^2 = a_0^2 + (H + h_0)^2 = H^2 + 2(R_E + H)h_0, \quad (2)$$

where d_0 is the distance from the first ($j = 0$) delay class to satellite and find delay to the satellite $\tau_0 = d_0/c$.

- STEP 3 : At j th iteration, determine the $(j + 1)$ th delay class. Select constant height $h_j = A/(2\pi R_E)$. Then, compute distance parameters

$$a_j^2 = R_E^2 - \left(R_E - \sum_{i=0}^j h_i \right)^2 \text{ and } d_j^2 = a_j^2 + \left(H + \sum_{i=0}^j h_i \right)^2 = H^2 + 2(R_E + H) \sum_{i=0}^j h_i. \quad (3)$$

and find delay of $(j + 1)$ th delay class to satellite, $\tau_j = d_j/c$.

- STEP 4 : Repeat this procedure for $j = 1, 2, \dots, \kappa - 1$.

Heights h_j , defining $(j + 1)$ th delay class are selected, so that the area of each zone, defined by heights h_{j-1} and h_j , is equal to that of the defined cup. The algorithm presupposed a given number of delay classes κ . If that is not given, it can be derived as

$$\kappa = \min \left\{ \omega : \cos^{-1} \left(1 - \frac{\sum_{i=0}^{\omega} h_i}{R_E} \right) \leq \theta_{min} \right\}. \quad (4)$$

The satellite footprint is sequentially covered with delay classes until angle θ_{min} is reached.

1.3.1 Delay class determination and resource allocation

We study the problem of determining the serving delay class for a call. When a call is initiated, resources are assigned to it after a call request message that contains current satellite and beam

identities and the propagation delay to satellite. If a beam or a delay class handover occurs, the new delay class can be determined, since timing to the current satellite is maintained. However, after a satellite handover, synchronization is lost, since satellite synchronization systems are independent. Although a satellite handover is less frequent than a beam handover, it can certainly occur since satellite footprints move fast on the earth surface. A satellite handover may occur during the call, if a call is either of long duration or is carried by an edge beam. The derivation of the delay class of the new satellite is important for reliable resource assignment and the delay from satellite-UT delay needs to be determined. Two methods can be used by the GW to determine this delay. According to method 1, the GW can retrieve the most recent estimate of UT position and associate it with the new satellite ephemeris data to find an estimate of the delay. In method 2, the GW can request an explicit measurement report from the UT. The first method is faster and easier to implement, while the second one is more accurate but also more time- and bandwidth-consuming. The first method should be given priority and used whenever the estimated UT position is accurate enough to provide a good estimate. A UT position error is acceptable if it does not lead to an incorrect delay class. If the UT position estimate is not accurate enough, the system should resort to the second method. The question that arises is when to use each method, so as to minimize the incurred signaling load. One could argue that Global Positioning System (GPS) would facilitate accurate computation of delay. However, in the case of a satellite handover, delay measurement with GPS would require message exchanges between UT and several satellites, which would consume power and bandwidth. In addition, delay determination methodology must maintain some compatibility with second generation technologies that do not include GPS in their standards. Hence, the delay can be determined with the two aforementioned methods.

1.3.2 UT position error tolerance region

Each UT is characterized by an actual location on earth and a delay to a satellite. Only estimates of these quantities are available and estimated UT position is referred to as known UT position. Consider a beam that is served by two delay classes (figure 5). Let the delay classes be defined by delays $T_{0,i} \pm \delta T_p$, for $i = 1, 2$, with $T_{0,1} < T_{0,2}$. When the UT is located in regions 1 or 2, it is assigned to a carrier of the corresponding delay class (1 or 2). Region 3 corresponds to delay range $[T_{0,2} - \delta T_p, T_{0,1} + \delta T_p]$ and is the overlap region of delay classes. Range δT_p determines the length of overlap region and is usually fixed for any pair of delay classes. While in the overlap region, a UT can be served by carriers of delay class 1 or 2. The best quality carrier of two carrier groups is selected. The UT may belong to one of the three depicted regions in a beam. However, because of delay estimation errors it may *seem* to reside in a different region from the actual one. As is explained in [2], an incorrect delay class assignment occurs if the difference between actual and known position corresponds to delay difference that is greater than the length of overlap region. Thus, the length of delay class overlap regions in the new beam after satellite handover must be determined. For a beam with n delay classes and $n - 1$ overlap regions, finding the minimum length overlap region and converting it to distance corresponds to computing the worst case error in position determination for which a correct delay class assignment is derived. The correct delay class must be assigned, so that UT is assigned to an appropriate carrier. The following methodology

is proposed to solve the problem.

- STEP A : Divide the set of satellite beams into subsets $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_L$, so that each beam in \mathcal{B}_j has κ_j delay classes. Beams of subset \mathcal{B}_i are type- i beams and form a toroid.
- STEP B : For each delay class pair $(i, i + 1)$ of beams in \mathcal{B}_j , find lengths of overlap regions x_{ij} , as $x_{ij} = T_{0,i} - T_{0,(i+1)} - 2|\delta T_p|$, where $T_{0,i} > T_{0,(i+1)}$. For subset \mathcal{B}_j with κ_j delay classes, $i = 1, \dots, \kappa_j - 1$. Then, for each \mathcal{B}_j , select $x_j = \min_i x_{ij}$, to account for minimum length overlap region for a beam with many delay classes. Then, x_j is the maximum allowed error between actual and estimated delay for beams in \mathcal{B}_j , for correct delay class assignment. Thus, $x_j = \max_k |T_{p_k} - \tilde{T}_{p_k}| = \hat{e}_{max,j}$, with T_{p_k}, \tilde{T}_{p_k} being actual/estimated delay for user k in \mathcal{B}_j .
- STEP C : Find distance tolerance $\Delta d_j = c \cdot \hat{e}_{max,j}$, $c = 3 \times 10^8$ m/sec is the light velocity.
- STEP D : For each subset \mathcal{B}_j and delay class $i = 1, \dots, \kappa_j$, compute the two distance extremes to satellite, $d_{i,j}^\pm = c \cdot T_{0,i} \pm \Delta d_j$ and corresponding central angles with the law of cosine [3]:

$$\theta_{i,j}^\pm = \frac{R_E^2 + (R_E + H)^2 - (d_{i,j}^\pm)^2}{2R_E(R_E + H)}, \quad (5)$$

- STEP E : Compute radius of each circular tolerance region $L_{i,j} = (R_E/2) |\theta_{i,j}^+ - \theta_{i,j}^-|$, and for the beam set with κ_j delay classes, select $L_j = \min_{i=1 \dots \kappa_j} L_{i,j}$ to account for worst-case error.

1.3.3 Delay class assignment

We describe the procedure to select the appropriate method (method 1 or 2) and assign the correct delay class. Delay class assignment is equivalent to estimation of UT position. The region in which the UT resides and the quality of position estimate determine the method to be used. We assume that delay and Doppler frequency offsets are available and the estimated UT position is derived from estimates of delay and Doppler frequency that are randomly distributed around the actual values, according to a Gaussian distribution. Let UT_{act} and UT_{known} be the actual and known UT positions and assume that the beam has $n > 1$ delay classes. The known UT position corresponds to non-overlap region k , $k = 1, \dots, n$, or to overlap region ℓ between delay classes $\ell, \ell + 1$, $\ell = 1, \dots, n - 1$. In the first case, method 1 can be used if UT position error is less than the minimum length overlap region. In the second case, we need to compute $\tau_1 = T_{0,(\ell+1)} + \delta T_p - \tilde{T}_p$ and $\tau_2 = \tilde{T}_p - (T_{0,\ell} - \delta T_p)$, which denote distances of the known position from the two sides of the overlap region. The procedure to make the delay class assignment is as follows.

- STEP 1 : Derive tolerances L_j for each beam subset \mathcal{B}_j .

- STEP 2 : For every UT in the satellite footprint, execute steps 2-5.
- STEP 2 : Compute actual propagation delay and Doppler frequency offset. Generate delay and Doppler frequency offsets, based on Gaussian distribution and derive known UT position.
- STEP 3 : Compute distance between estimated and actual UT position and determine the beam and beam type \mathcal{B}_m of known UT position.
- STEP 4A : Known delay \tilde{T}_p corresponds to a UT position in a non-overlap region ℓ . Then, if $|UT_{act} - UT_{known}| < L_m$, then use method 1 to find delay, else use method 2.
- STEP 4B: Known delay \tilde{T}_p corresponds to UT position in an overlap region ℓ . Then, if $|UT_{act} - UT_{known}| < \min\{\tau_1, \tau_2\}$, then use method 1 to find delay, else use method 2.
- STEP 5 : Use this delay value to assign the call to a delay class.

If parameters τ_1, τ_2 are computed, $\min\{\tau_1, \tau_2\}$ is the closest distance from the overlap region boundary. The condition $|UT_{act} - UT_{known}| \geq \min\{\tau_1, \tau_2\}$ means that the error region covers overlap region ℓ and part of non-overlap regions ℓ or $\ell + 1$. Then, only method 2 can guarantee a correct delay class assignment. Otherwise, if $|UT_{act} - UT_{known}| < \min\{\tau_1, \tau_2\}$, the error region lies entirely in overlap region ℓ and method 1 can be used.

1.4 Handover issues in mobile satellite networks

The problem of resource allocation in mobile satellite networks is inherently related to handover [4]. Each time a handover occurs, a resource allocation problem needs to be solved and new resources need to be assigned to the user. The objective is to accommodate as many users as possible in the system and guarantee acceptable link quality for each user. Then, the system can respond better to a potential sudden load increase or link quality deterioration and the likelihood of blocking a connection request is minimized [5]. Efficient resource allocation and handover techniques aim at satisfying quality of service (QoS) requirements of users. The contributions of this research are (i) the establishment of a framework for reliable handover prediction at the GW, (ii) the consideration of meaningful criteria for selection of the appropriate satellite and beam, whenever the user can choose among several satellites and beams for handover and (iii) the proposition of methodologies with which handover signaling load is reduced. Several alternatives for handover prediction and path selection are proposed and evaluated [6].

1.4.1 Beam selection for measurements

The standard procedure of beam signal level monitoring in GSM terrestrial cellular networks is used. Each user monitors the received signal strength of a set of broadcast common channels

(BCCH), where each BCCH corresponds to an adjacent beam of a satellite. The GW or satellite base station (SBS) periodically commands the user to measure signals of all visible serving and non-serving satellites and creates a list of beams to be measured. These measurements will facilitate handover decisions. The list comprises the set of beams of different visible satellites that cover the user location, as well as a set of *approaching* beams of serving satellites. There can be one or two serving satellites, depending on the use of diversity. The beams of the first set are candidates for satellite handover, while the beams in the second set are candidates for beam handover. The procedure takes place at call origination and during the session of a call and is depicted in figure 6.

The problem of high handover signaling load becomes important in non-geostationary satellite systems, which are characterized by small beams and larger number of handovers. In order to efficiently utilize scarce bandwidth resources, signaling load should be reduced. The required time to perform measurements and ultimately make the handover decision depends on the size of the measurement list, which increases with the number of visible satellites and with diversity. During connection setup, the objective is to minimize this time, so that transition to traffic mode occurs fast. An alternative method in which only one beam from each visible satellite is measured requires less time to accomplish. Although such a method is desirable during connection setup, it may result in a suboptimal path during traffic mode. Numerical results in [6] indicate a significant improvement in measurement time and associated signaling load with this technique.

1.4.2 Path selection

Path selection algorithm is an inherent component of resource allocation and takes place after beam selection and before handover. Each entry k of the list is a pair of satellite and beam indices (s_k, b_k) . The list is modified as follows. All possible combinations (k, ℓ) of single elements are created and appended to the list. Entries with a measurement below a given threshold are eliminated since they indicate unreliable connection. The list is ranked based on a preference factor, so that the first entry has the highest preference factor. Each entry is a single or a diversity "path" that is eligible for resource allocation. For entry k the preference factor P_k is a function of satellite elevation angle θ and the difference of the received signal strength from threshold. If the entry is a diversity path, the preference factor depends also on the azimuth separation angle ϕ between satellites. A path with high satellite elevation angle is more preferable for assignment, since the likelihood of call blocking due to some obstacle is reduced. Furthermore, a high signal strength provides better quality and reduces the probability of blocking due to unacceptable signal level. Finally, a big azimuth separation angle is more preferable, since there are fewer chances that both paths get corrupted due to an unpredictable blockage. Thus for the more general case of a diversity path,

$$P_k = A \left(\frac{\theta_{k,1}}{\pi} + \frac{\theta_{k,2}}{\pi} + \frac{\phi_{k,12}}{2\pi} \right) + B (I_{k,1} + I_{k,2}) \quad (6)$$

where A, B are factors that capture the relative significance of these criteria.

1.4.3 Handover prediction algorithms

When a UT enters a beam, it is mapped to a coordinate system [7], until a time interval of some length (e.g., 1 min) is found, so that the UT resides in the current beam at the beginning of the interval and in a different beam at the end of the interval (figure 7). Handover must occur sometime during this interval and predicted handover time is the mid-point of the interval. Motivated by the fact that additional handovers lead to increased signaling load and delay, we present a new criterion for handover. The key idea is that the user should be forced to reside in a cell as much as possible, so that handover rates are minimized. Upon creation of the list of candidate paths for transition, no preference factor computation is required. Instead, the entry with the beam in which the mobile is predicted to stay the longest is selected for handover. This beam may or may not belong to the serving satellite. This criterion involves fewer computations and measurements than path selection. User residence time for each beam in the list is computed by standard coordinate system computations without the need for elevation angle computations. More importantly, no signal strength measurement exchange is necessary.

This prediction technique assumes stationary UT. If the UT is in motion, successive handover predictions may lead to prediction time errors. Beams move at different rates relative to the earth surface, since they follow different paths around the earth. Edge beams move at lower rates than central ones. UT mobility has significant effect on beam transitions in the slow-moving beams and small effect on fast-moving ones. Additional results on the effect of mobility on anticipated handover times can be found in [7].

For the satellite handover, the serving satellites at current time t_c must be identified and updated to a future time t_f with satellite ephemeris data. If $\theta_{UT}(t_f) \leq 10^\circ$ or $\theta_{GW}(t_f) \leq 10^\circ$, then a handover occurred at some time $t^* \in (t_c, t_f)$. We apply bisection on that interval and stop after n^* iterations, when $t_{f,n^*} - t_{c,n^*} < W$. The predicted satellite handover time is $t_s = (t_{f,n^*} + t_{c,n^*})/2$ [6].

1.5 Numerical results

A MEO mobile satellite system with the specifications of ICO satellite system was simulated. A satellite has 163 beams of 8 types. A UT can be served by a satellite when its elevation angle θ exceeds $\theta_{min} = 10^\circ$, which is typical for a rural environment. Within a satellite footprint, delay class positions are computed and concentric rings as those in figure 2 are defined. In order to serve beams close to nadir with one delay class, we set the closest delay class to be the ring at $\theta = 60$ degrees. By applying the delay class determination algorithm, we found that the maximum number of delay classes that ensures beam coverage is 3 and beams are covered by 1,2 or 3 delay classes. The distance tolerance values that guarantee correct delay class assignment for beams with 2 and 3 delay classes were found to be $142km$ and $24km$ respectively. Position determination is more sensitive in beams with 3 delay classes, due to delay class proximity and small length of overlap

regions. UTs are covered by one GW. Satellite, beam and delay class handovers can occur for UTs. In satellite handover, a UT always uses one satellite as long as it remains visible. A beam handover occurs when a UT moves in the coverage area of another beam within the satellite. A delay class handover occurs when the UT moves to another delay class within a beam. Since beams or delay classes overlap, a beam or delay class handover occurs randomly during the time when the UT is in the beam or delay class overlap area. We simulated an one-hour satellite revolution. Calls in different beams arrive in Poisson streams of rate λ and their durations are exponentially distributed with mean $\tau = 1/\mu = 150\text{sec}$, which is typical for voice transactions. Traffic intensity is measured in Erlangs (E), $E = \lambda\tau/60$. Although handovers are affected more by satellite velocity, a simplistic UT mobility model is also used with a random number of UTs having speed uniformly distributed in $[0, 72]$ km/h and direction in $[0, 2\pi]$.

The carrier grouping method leads to efficient resource management and reduced blocking rate. The presented method does not significantly affect blocking rate in moderate size beams with one delay class, but it offers a clear advantage in edge beams with large delays. We performed experiments for an edge beam and measured performance for randomly generated UTs within the beam region, assuming that carriers are allocated to the beam with FCA. In the scenario with no carrier grouping, UTs were allocated randomly in carriers irrespective of their location in the beam. Thus, UTs assigned at the same carrier had different Tx/Rx offsets so as to maintain orthogonality. In the second scenario, the beam was divided in 3 delay classes, the set of carriers was divided in three subsets and each subset was mapped to a delay class. Calls belonging to a delay class were assigned to a carrier of that class. Each time a carrier of a delay class became full, UTs were assigned to another carrier of that delay class. Overlap area calls were assigned to the least loaded carrier of those in the delay classes. When all carriers of a delay class were full, the call was blocked. The blocking probability is the ratio of blocked calls over total number of calls. Experiments with 12 and 45 carriers were performed and results are illustrated in figure 8. The improvement in performance due to carrier grouping is evident. For example, in the case of 12 carriers and a traffic of 25 Erlangs, blocking probability was reduced by almost 50%. For higher loads or for a larger resource pool of 45 carriers, the advantage of carrier grouping becomes more substantial.

We also measured the resulting satellite, beam and delay class handover rates in our simulation. In figure 9, the handover rates are depicted as a function of time, for a square of $15^\circ \times 15^\circ$ within 15° longitude and latitude distance from the GW. The aggregate handover rates for all traffic carried by the GW is also depicted. A first observation is that all handover rates converge to a steady state after a transition interval of about 2,500 seconds. In steady state, about 30 delay class transitions per minute occur. Delay class handover rate depends on whether the UT is located in edge beams with more than one delay classes. The ratio of satellite, beam and delay class handover rates in steady state was is 2 : 7 : 3. For locations further from the GW, this ratio was 1 : 15 : 1. Furthermore, delay class handover rate becomes in a way predictable, since it occurs only within beams with more than one delay class.

In figure 10, we present comparative results about satellite and beam handover rates by using UT position or maximum beam residence criterion in a region with moderate load (0.92 calls/sec).

With the second criterion, a reduction to beam handover rate up to 85 – 90% was observed, while for heavier traffic load this reduction reached 35% – 40%. It is important to observe the increased satellite handover rate for some time periods. However, the effect of increased satellite handover rate is low, since the satellite handover rate at steady state is as low as 3 – 4 handovers per minute.

2 Information Dissemination in Satellite Networks

2.1 Push-Based Multi-stage Information Dissemination

Broadcast data delivery is an effective way to disseminate information to massive user populations. The two foremost advantages of broadcast data delivery are scalability to large user populations, and high availability to geographically distributed mobile users. However, users need expensive and cumbersome equipment to receive and transmit satellite signals. Furthermore, as the amount of information being broadcast increases average user latency increases as well. Often users in a geographical locality have similar interests, which can be better served by employing a local broadcast schedule. In this context, a two stage satellite-terrestrial wireless broadcast system can provide more efficient service in terms of lower average user latency, and cheaper and more convenient user equipment. In such a system, main server broadcasts information via satellite to the geographically distributed local ground stations. Every ground station has limited buffer capacity to store the data broadcast by the satellite. According to their buffer content and the interests of their users local stations deliver information via terrestrial wireless channel.

Main contribution of this research initiative is the identification and treatment of the issues that arise in novel multi-stage broadcast scheduling-caching applications. Specifically in our work, we investigate the design of the satellite schedule, the cache management policies at the ground stations and the terrestrial schedules that arise in joint fashion in the context of our problem. Our results demonstrate that such a two stage system is feasible and can provide more efficient data delivery compared to a single stage system.

2.1.1 System Model

Under the *push-based* broadcasting approach as depicted in Figure 11, a server continuously and repetitively broadcasts data to a user community without any feedback about the users' needs due to the limited uplink communication capability from users to the server. The data is delivered in terms of *items*, where an item is a fixed sized information unit. The items are broadcast in a particular order that is called a *schedule*. Time is slotted, and duration of each slot is equal to the transmission time of an item on the broadcast channel. When a user needs a certain item, it monitors the broadcast channel until the desired item is detected. There is some latency for the fulfillment of the needs of a user, since the server can only broadcast a single item per broadcast channel at any time instant. This latency depends on the broadcast schedule, as well as the user

request pattern.

Figure 12 depicts the two-stage satellite-terrestrial system that is going to be discussed in the remainder of this paper. In this setup, the main server transmits information items via satellite to the geographically distributed local ground stations with a transmission rate of R_s items per unit time. The ground stations have a limited local buffer capacity of C items. Each ground station retransmits the items stored in the local cache to a local user community with transmission rate R_g items per unit time, where $R_s = r \cdot R_g$, and r is a positive integer. We investigate the problem of minimizing the overall mean user latency by developing intelligent satellite and terrestrial broadcast schedule and local server cache management policies. In this context, the terrestrial schedule should try to match the user access characteristics as much as possible, given the satellite schedule and the available items saved in the cache. The cache management policy should in turn store the items broadcast by the satellite that are expected to be broadcast by the terrestrial schedule. The satellite schedule should deliver the items in the order that the underlying cache management and terrestrial scheduling policies can minimize the local average user latency.

2.1.2 Joint Cache Management and Terrestrial Broadcast Scheduling

Even if we use naïve cache management and terrestrial scheduling techniques, the resulting performance improvement with the two-stage systems is significant. Let's illustrate this with the following example. Assume that there are four groups of users that are interested in only four items. Each item is requested with the same probability. Neither of the items requested by the users of different groups are the same. Therefore, the main server broadcasts sixteen items to the local ground stations of these four groups. Assume that the satellite schedule is uniform, that is, these sixteen items are broadcast sequentially and periodically. Each local ground station can store only two items at any given time instant and the satellite channel is twice as fast as the terrestrial channel. An intuitive and simple cache management procedure is to "store any new item of interest received from the satellite schedule and remove the item that is the oldest in the cache." Under this policy, the cache content evolves as depicted in Figure 13. An equally simple terrestrial scheduling policy "broadcasts at a given instant the item residing in the cache that is broadcast least recently." Under this policy the terrestrial schedule evolves as depicted in Figure 13. It is easy to see that the resulting terrestrial schedule has a mean response time of 4.25 satellite slots. Instead of two stage broadcast data delivery if a single-stage broadcast scheduling policy was deployed, the single-stage broadcast schedule would have to service all four user groups and the best schedule would then be the uniform satellite schedule that is also depicted in Figure 13. The mean response time of this schedule is 8 satellite slots. The improvement in the average user latency by employing the two-stage system is almost 50%.

If all items that are requested by the users in a geographical locality are available at the local ground station, then there is no need for cache management, and the problem reduces to single stage broadcast scheduling problem, which is extensively studied. In these studies, it has been noticed that in any broadcast scheduling system, there are two quantities related to each item i that affect

the scheduling decision at slot n . These are, elapsed time $w_i(n)$ since the last transmission of item i and λ_i the request generation rate for item i . The likelihood of item i being transmitted at n increases with λ_i and $w_i(n)$. In [10], Su, Tassiulas and Tsotras considered a set of policies for a single-stage system, where the broadcast scheduling is determined based on priority indices of the items. The priority index of item i is the product $\lambda_i^\gamma w_i(n)$, where γ is an exponent that reflects relative importance of λ_i versus $w_i(n)$. It is shown that $\gamma = 0.5$ gives the best performance among these priority index policies. $\gamma = 0.5$ refers to a special case, where the index is $\sqrt{\lambda_i w_i^2(n)}$. Note that $\frac{1}{2}\lambda_i w_i^2(n)$ is the *aggregate expected delay* experienced by item i requests since the last transmission of item i before slot n . Su et al., showed that this, so called *Mean Aggregate Delay (MAD)*, algorithm gives performances close to the lower bound.

The single stage and the proposed two-stage systems differ in the availability of items for scheduling at a given slot. In two-stage system, only the items that are stored in the cache of the local ground station are available for scheduling. Intuitively, we want the terrestrial schedule of a local station in the two-stage system be same as the schedule in a single-stage system that is serving exclusively the local user group. In order to achieve this goal we propose to use the MAD scheduling policy as the local terrestrial scheduling policy and develop good cache management policies that work on top of the MAD scheduling policy. A good cache management policy should try to feed the scheduler with the items that are going to be broadcast in the coming slots according to the MAD policy as if all items are available to the scheduler. We now describe two cache management algorithms that comply to this intuition.

TWT Algorithm

Due to the lack of uplink request channel, the only information about the user accesses that the local stations can determine at a given slot, is the expected number of users waiting for service. The expected user backlog for item i observed by an arbitrary local station at satellite slot n is defined as $\bar{X}_i(n) = \lambda_i w_i(n)$. The evolution of the expected user backlog for item i at a local station is depicted in Figure 14. The aggregate average delay of item i at slot k is defined as $\frac{1}{2}\lambda_i(w_i(k))^2$, and it is the striped area between t_0 and k in Figure 14.

Assume that at satellite slot k item i is broadcast by the satellite. The local station has two options; to fetch item i and replace another item in the cache, or to leave the cache content unchanged. If item i is not fetched at k , the earliest time it can be scheduled at the terrestrial schedule is at satellite slot $k + l_i(k)$, i.e. when item i is re-transmitted by the satellite. At that instant the expected item i backlog in terms of satellite slots will be given by, $b_i(k) = \bar{X}_i(k + l_i(k)) = \frac{1}{2}\lambda_i[w_i(k) + l_i(k)]^2$. Observe that $b_i(k)$ refers to the shaded area in Figure 14. The mean aggregate delay depends on the evolution of the expected backlog of items, $b_i(k)$. Thus, higher is the $b_i(k)$ higher is the mean aggregate delay. TWT algorithm fetches an item, if the total waiting time of that item at its next satellite re-transmission is high.

ETT Algorithm

One weakness of TWT algorithm is that it is a *myopic* algorithm in the sense that it does not take into consideration the future desired terrestrial schedule. Unlike TWT algorithm, ETT algorithm first estimates the future evolution of the terrestrial schedule at an arbitrary satellite slot k . According to this estimate, ETT algorithm then checks whether item $i \in B(k)$ is broadcast or not at terrestrial channel prior to item i 's first re-transmission at the satellite channel after the k th slot. If it is not broadcast, then there is no need to keep item i in the cache. Else, ETT algorithm calculates the reward of keeping item i in the cache at slot k . Let t_1 correspond to the satellite slot the terrestrial transmission of item i is predicted to take place, and let t_2 correspond to the satellite slot the first satellite re-transmission of item i after k th slot takes place (see Figure 15). If item i is not stored at k , the earliest time item i can be scheduled at the terrestrial schedule is at t_2 . Item i is not supposed to be broadcast at the terrestrial channel until t_1 , so in fact there is no benefit of keeping item i stored in the cache until t_1 . Thus, the reward of keeping item i in the cache is only the additional waiting time from t_1 to t_2 that is eliminated by the broadcast of item i at t_1 . This reward is illustrated as $d_i(k)$ in Figure 15. Hence, ETT algorithm stores an item i at slot k , if the reward of keeping item i , $d_i(k)$, is sufficiently large.

In Figure 16 we demonstrate the performance of TWT and ETT with respect to a lower bound that we derived in [8]. In this example, the probability that item i is requested is given by Zipf 1 distribution, i.e. $Pr(i) = \frac{c}{i}$, where c is a normalization constant [11]. ETT outperforms TWT with quite a margin. This is because ETT estimates the future terrestrial schedules and tries to buffer items from satellite schedule to satisfy this estimated schedule. More results can be found in [8].

2.1.3 Satellite Schedule Design

In general, the users in different local stations have different interests, i. e. different access probability distributions. Therefore a satellite schedule, that is suitable for a user group, may result in the worst possible performance for another. Joint design of the satellite schedule, cache management strategies, and terrestrial schedules may provide a better overall performance. We assume that one of the proposed cache management algorithms, and the MAD scheduling policy are employed together at the local stations. In particular, we assume that the ETT algorithm is used as the cache management algorithm. We consider off-line design of the satellite schedule. As a heuristic, we consider the case where the satellite schedule is determined by the MAD policy. According to this problem setup the only question that remains to be answered is, which access probability distribution should be used to determine the satellite schedule that minimizes the overall average response time of the users. We consider two cases: aggregate probability distribution where all user groups have equal weight, and WSAT where the weight of each user group is inversely proportional to its local buffer size. Following table demonstrates the performance of these approaches on an example.

Group	Distribution	Cache size	MRT	
			WSAT	Aggregate
1	z1(1,20,100)	$C_1 = 5$	24.48	28.31
2	z1(20,20,100)	$C_2 = 5$	24.83	28.30
3	z1(40,20,100)	$C_3 = 5$	25.64	30.20
4	z1(60,20,100)	$C_4 = 10$	24.65	21.72
5	z1(80,20,100)	$C_5 = 20$	8.07	8.07
Overall			21.534	23.32

Table 1: Example

2.2 Efficient Resource Allocation for Pull-Based Information Distribution

As it is the case in the Internet, users often *pull* information rather than waiting for the required information to be *pushed* to them. As discussed previously, many users lack a high speed network connection. Users usually have a connection to a local network, where a local gateway station is connected to the backbone network. In rural areas, satellite link appears to be a feasible method of connection for these gateways. However, satellite bandwidth is expensive and it will be cost-effective to serve the requests locally. Thus, local gateways usually employ caches to store the popular objects. Until now, system architecture is similar to the system discussed in the previous section. However, we further consider the case where local gateways that are in close proximity to each other are connected via a terrestrial link so that the user requests that cannot be satisfied at one gateway are directed to another gateway. Such a flexibility to forward user requests among the gateways allows user requests to be satisfied locally and most probably with much lower latency. The architecture is summarized in Figure 17.

We consider a business model in which the primary customers of the local gateways are the publishers (origin servers). The objective of the publisher is to maximize its net benefit. The net benefit is the utility a publisher receives by delivering its content to the users with delay \mathcal{D} , minus the cost of the resources required to ensure that delay. A cost arises, because the gateway has limited resources such as bandwidth, cache, etc. and uses prices to induce efficient allocation. The objective of a gateway is to maximize its revenue, while satisfying the requested QoS. Publishers act non-cooperatively and try to maximize their own benefit regardless of the rest of the servers.

Content Delivery problem can be separated into two: content distribution and user request routing sub-problems. Given user request arrival rates to the surrogates, the distribution sub-problem determines the best dissemination strategy, i.e. the set of objects disseminated to each surrogate, maximizing its net benefit. Meanwhile, given the surrogate cache content, the request routing sub-problem determines the best user request-routing strategy, i.e. the set of user requests served by each surrogate, again maximizing the net benefit of the publishers. Although these problems are jointly related, for the purpose of designing a practical algorithm, we envision an iterative scheme, in which each sub-problem is solved separately. In this scheme, the distribution sub-problem determines the dissemination strategy according to the arrival rates as determined by

the request-routing sub-problem. The request routing sub-problem determines the routing strategy according to the cache content as determined by the distribution sub-problem. These sub-problems update their decisions iteratively according to the output of the other. We expect that the solution to the overall content delivery problem is (near-) optimal, if the solutions to these two sub-problems are (near-) optimal. The formulation and the aim of these two problems closely resemble each other. In the following, we describe the results for the distribution sub-problem with the understanding that similar results apply to the routing sub-problem as well. For further details on the routing sub-problem we refer the readers to [12].

2.2.1 System Model

The objective of the distribution sub-problem is to maximize the reduction in average user latency by the use of surrogate caches, while the total cost of service is less than the highest investments the publishers are willing to make for the caching resources. Figure 18 illustrates the network set-up that we are interested in this section. There are several LANs from which the user requests originate. Every user is interested in one or more of the objects. The user requests are first intercepted by the request-routing sub-system. The request-routing subsystem then forwards the user request to an appropriate surrogate, perhaps the one with the highest probability of delivering the requested object with minimum latency. For the distribution sub-problem, we assume the routing decision is taken, and the user request arrival rates to each surrogate are determined. Surrogates are located between the user networks and the origin servers. Thus, users are always at most two hops¹ away from the content. A user request is first checked at the corresponding surrogate. If the requested object is available at the surrogate, request is immediately served. Otherwise, request is forwarded to the web server, where the object originally resides.

Surrogates have limited cache size, and the cache is shared among the publishers. Surrogates charge the publishers for the portion of the cache their content occupies. We assume that the surrogates does not cooperate – other than announcing their resource prices – probably due to high messaging and processing overheads associated with cooperative caching. Instead, they act non-cooperatively with the objective of maximizing their individual revenues. Surrogates compete among each other to store the publishers content. Publishers do not collaborate either and they try to purchase as much cache space as possible with the objective of maximizing their net benefit. Assume that there are I different publishers (origin servers) and J surrogates present in the network. User requests arrive from N different user LANs. Let $\lambda_i^{j,n}$ denote the total request arrival rate from LAN n at surrogate j for the content in the i th publisher. Let $\lambda_i^j = \sum_n \lambda_i^{j,n}$ be the total arrival rate to the surrogate j for the content in publisher i . The user interest in the objects of the publishers is distributed according to Zipf distribution [11].

Let B_i^j be the investment of the i th publisher in the j th surrogate. Let $\mathcal{B}_i = \sum_j B_i^j$ be the

¹Notice that user *packets* may traverse multiple network elements/links to reach the server.

total investment of the i th publisher. It is assumed that the information stored in the servers is continuous and can be replicated continuously to a surrogate. Total information available at the publisher i is χ_i . Publisher replicates its most popular part of the content to the surrogates so that the cache hit probability is maximized. Assuming that C_i units of cache space is allocated to the publisher, the probability that an incoming user request is satisfied at the surrogate is given by,

$$Pr(\text{hit}|C_i) = \int_0^{C_i} q(x) dx = \int_0^{C_i} \frac{1-\alpha_i}{\chi_i^{1-\alpha_i} x^{\alpha_i}} dx = \left(\frac{C_i}{\chi_i}\right)^{1-\alpha_i}.$$

Let p_j denote the price of the unit cache space in surrogate j . Let the pricing policy, $\mathbf{p} = (p_1, p_2, \dots, p_J)$, denote the set of unit cache space prices of all the surrogates in the network. Let d_{ij} denote the additional average delay a user request forwarded from surrogate j to the origin server of publisher i will experience. Note that if we consider geosynchronous earth orbit satellites, we can assume that d_{ij} is approximately the same for all i, j . However, in this study we consider a more general system, where a low-earth orbit satellite network (LEO) may be used to connect the origin servers and the gateways. LEO satellites have propagation delays comparable to the terrestrial networks and thus the locations of the origin servers and the gateways have effect on the observed delays.

2.2.2 . Optimal Publisher and Surrogate Strategies

We first consider the optimal strategy for the publishers. Let x_i^j be the cache space allocated to publisher i in surrogate j . If i th publisher's investment in the j th surrogate is B_i^j , then the total cache space allocated to the content of publisher i in surrogate j is $x_i^j = \frac{B_i^j}{p_j}$. The average reduction in the user delay or equivalently the average net benefit that publisher i generates by B_i^j investment in surrogate j is $\lambda_i^j w_i(d_{ij}) \left(\frac{B_i^j}{p_j \chi_i}\right)^{1-\alpha_i}$. Define $\beta_i^j = \lambda_i^j w_i(d_{ij}) / \chi_i^{1-\alpha_i}$ as the gain factor for publisher i from surrogate j . The utility function, i.e. the total *additional* average benefit, $U_i(x_i)$, of publisher i is $U_i(x_i) = \sum_{j=1}^J \beta_i^j (x_i^j)^{1-\alpha_i}$. For a given pricing policy \mathbf{p} the publisher optimization problem (S) can be given as:

$$(S) \quad \max_{\{x_i^j\}_{j=1}^J} \quad U_i(x_i)$$

$$\text{subject to} \quad \sum_{j=1}^J x_i^j p_j \leq B_i.$$

Since $U_i(x_i)$ is a concave function, and the constraint set is compact, there exists a unique solution to (S).

Lemma 1 $x_i^{j*} = \frac{\left(\frac{\beta_i^j}{p_j}\right)^{1/\alpha_i} B_i}{\sum_{k=1}^J p_k \frac{\beta_i^k}{p_k}^{1/\alpha_i}}$ is the unique optimal solution to the optimization problem (S).

We now consider the optimal pricing strategies of the surrogates maximizing their revenues. Let $\mathbf{p}^{-j} = (p_1, p_2, \dots, p_{j-1}, p_{j+1}, \dots, p_J)$ be the set of unit cache space prices of all the surrogates in the network except the j th one. We assume that there is no collaboration among the surrogates, and each surrogate tries to maximize its revenue non-cooperatively.

Lemma 2 Surrogate revenue $r_j(p_j) = \sum_i x_i^j(p_j) p_j$ is maximized under a given fixed pricing policy \mathbf{p}^{-j} , when p_j satisfies $\sum_{i=1}^I x_i^j(p_j) = C_j$, i.e. when the total publisher demand is equal to the surrogate cache space.

Until now, we discussed the optimal strategies of the publishers and the surrogates given that system is at a steady state. However, we have not discussed whether such a steady state exists. Notice that when a surrogate re-evaluates its pricing policy according to the pricing policies of the rival surrogates, the remaining surrogates will do the same. For each different pricing policy the publishers' optimal investments will be different as well.

In order to understand the behavior of the surrogates, we model the two-stage publisher-surrogate system as a non-cooperative game. In this *publisher-surrogate distribution game*, $\Gamma(J, S, P)$, the players, J , are the surrogates, the strategy set S_j for a surrogate j is given by the surrogate's unit cache space price and the payoff function $P_j(s)$ of each surrogate j is given by the profit of the j th surrogate. This system is similar to the Cournot oligopoly discussed in the economics literature. Assume that each surrogate has a fixed cost for its cache, but has no control over the size of the cache, i.e. the size of the cache is determined prior to the system operation. We first show that this game has a Nash equilibrium solution, where no surrogate has incentive to change unilaterally its strategy, since each surrogate maximizes its own individual payoff given the strategies of others.

Theorem 1 *The non-cooperative publisher-surrogate distribution game $\Gamma(J, S, P)$ has at least one Nash Equilibrium solution.*

Pareto optimality is the relevant criteria in a multi-objective problem setting. At the pareto optimum solution, one can find no other feasible solution that increases some objectives while not decreasing at least another objective. We have shown in [9] that if the equilibrium is unique, then the outcome of the non-cooperative game is the optimal solution to the individual revenue maximization problems of the publishers. If there are multiple equilibria, however; the resulting cache allocations, \mathbf{x} , are only locally optimum. Unfortunately, there are often multiple Nash equilibria and depending on the initial prices as well as price update strategies the outcome of the game may not always be

the optimal solution. In [9], we discuss a special case of the surrogate cache allocation problem, where the delay between each publisher-surrogate pair is the same and user request arrival rates to each surrogate and Zipf distributions for each publisher are identical. For this case, we determine the condition for which unique equilibrium exists.

2.2.3 Selection of Publisher Investments

Until now, we assumed that publisher investments are fixed and known. We now discuss several possibilities for determining the investment amounts. Consider the following optimization problem (P). The objective of (P) is to maximize the total publisher utility given the limited cache sizes. We define the solution of (P) as the *system optimum* solution.

$$(P) \quad \max_{\{x_i^j\}_{i,j}} \sum_i U_i(\mathbf{x}_i)$$

$$\text{subject to } \sum_i x_i^j \leq C_j, \quad j = 1, 2, \dots, J.$$

This problem can be solved by considering the *dual* problem and using the gradient projection method. One important feature of this method is the fact that it can be implemented in a distributed fashion. The solution of this problem dictates the publishers to pay the cost associated with the *system optimum* solution. Notice that a server may be individually better off by investing less than the amount dictated by the system optimum solution. Thus, a system optimum solution can only be achieved by the distributed publisher-surrogate game by enforcing the publishers to pay the charges associated with the system optimum solution. Without such enforcement, we encounter a different optimization problem where the investment amount is also a design parameter. In the resulting non-cooperative game, publishers select their investment amounts to maximize their *profits*.

Consider a publisher-surrogate game where a publisher selects an investment maximizing the "net" publisher benefit. That is, the difference of the additional average benefit generated from using surrogates and the cost of using them is maximized. Let T be the duration of time a publisher can rent the caching space from a surrogate. Then for a given set of surrogate prices \mathbf{p} the publisher optimization problem can be given as,

$$\max_{\{x_i^j\}, \mathcal{B}_i} U_i(\mathbf{x}_i) - \frac{\sum_j x_i^j p_j}{T} \quad (7)$$

$$\text{subject to } \sum_j x_i^j p_j \leq \mathcal{B}_i.$$

We developed a suboptimal investment strategy for the publishers by assuming that the change in the equilibrium prices is *small* when the change in the investments is small as well. For a given

investment B_i , the optimal cache allocation is given by $x_{ij}(B_i) = \frac{\beta_{ij}^2/p_j^2}{\sum_k \beta_{ik}^2/p_k} B_i$. Assume that p_j are the set of prices at the equilibrium. As for the key assumption of our derivation, we assume that for a small change in investment amounts, the change in the equilibrium prices is small as well. Specifically, in the vicinity of the set of investment amounts B_i , $\sum_k \beta_{ik}^2/p_k$ can be considered as *constant*. This assumption is verified through numerical studies. Assuming that this condition is satisfied, the optimal investment amount maximizing the net benefit is calculated.

2.2.4 Numerical Results

Figures 19 and 20 demonstrate the performances of the suggested approaches. Let $\zeta_i^j = w_i(d_{ij})/\chi_i^{1-\alpha_i}$. From the definition of ζ_i^j and assuming that $w_i(\cdot)$ is the same for all publishers, one can notice that by varying ζ_i^j , we basically change the delay between surrogate j and publisher i . As illustrated in Figure 19, as the skewness of the system increases the performance of game-theoretic algorithm gets better compared to the conventional algorithm. However, this improvement is larger for low values of ζ_i^j . This means that surrogates need not be very close to the users. We may place surrogates at locations at somewhat greater distance and thus service more users with less number of surrogates with sufficiently good system performance.

In Figure 19 the a priori fixed investments are $B_1 = B_2 = 10$. The system optimum solution maximizes the total publisher utilities. As expected we observe that the total utility achieved by the investments associated with the system optimum solution is indeed higher than the fixed investments. It is interesting to note however, while the system optimum investment maximizes the *total* of publisher utilities, it reduces some of the publishers *individual* utilities. Furthermore, the investment amounts leading to the system optimum are usually very high, reducing any benefit received from the use of the caches in the first place. A short-sided objective of maximizing total system utility is thus seems not implementable, when the system is formed by selfish non-cooperative agents.

In Figure 20, we investigate the performance of our sub-optimal investment algorithm that prevents the abovementioned problem. We compare our results with the game theoretical algorithm with fixed investments, where the investments are 10 units. The system setup is similar to the previous numerical examples, and the improvement is considered for varying gain factor. We also depicted the reduction in the total publisher utility as compared to the optimal system utility. The cache rental period is 100 time units. In this example, we observe that the improvement in the publisher profits is significant when the publishers are allowed to optimally vary their investments. Furthermore, the decrease in the system efficiency is small as compared to the increase in the profit.

2.2.5 An Optimization Based Resource Allocation Algorithm

The results given in the previous sections suggest that we may use a price-directed market-based distributed algorithm for solving the two-stage publisher-surrogate cache resource allocation problem. We consider the following algorithm for this purpose: An initial set of prices is announced to the publishers. The publishers determine their resource (cache) demands according to these prices as well as the request rates, and the observed delays from the surrogates. The publishers request these resources from the surrogates. Prices are then iteratively changed to accommodate the demands for resources until the total demand equals to the total amount of resource available.

In [9], we show that this algorithm converges to the Nash equilibrium solution.

References

- [1] I. Koutsopoulos and L. Tassiulas, "A synchronization-based scheme for simultaneous full- and half-duplex communication in GSM-based MEO mobile satellite networks", *Proceedings of IEEE Globecom*, pp.276-280, 1999.
- [2] I. Koutsopoulos and L. Tassiulas, "Efficient resource utilization through carrier grouping for half-duplex communication in GSM-based MEO mobile satellite networks", *IEEE Transactions on Wireless Communications*, vol.1, no.2, pp.342-352, April 2002.
- [3] D. Roddy, "Satellite Communications", Mc Graw-Hill, 1996.
- [4] S. Tekinay and B. Jabbari, "Handover and channel assignment in mobile cellular networks", *IEEE Communications Magazine*, vol.29, no.11, pp.42-46, Nov. 1991.
- [5] L. Tassiulas, "Efficient spectrum management for high capacity wireless networks", *Proceedings of 8th IEEE workshop on Computer Communications*, pp.0.211-0.218, 1993.
- [6] I. Koutsopoulos and L. Tassiulas, "A unified framework for handover prediction and resource allocation in non-geostationary mobile satellite networks", *Proceedings IEEE Vehicular Technology Conference Fall*, vol.4, pp.2106-2110, 1999.
- [7] I. Koutsopoulos and L. Tassiulas, "Reliable handover prediction and resource allocation in MEO mobile satellite networks", *Proceedings of IEEE/IMACS Circuits, Systems, Communications Conference*, 1999.
- [8] O. Ercetin and L. Tassiulas, "Push Based Information Delivery in Two Stage Satellite-Terrestrial Wireless Systems," to appear in *IEEE Trans. on Computers*.
- [9] O. Ercetin and L. Tassiulas, "Optimal Competitive Resource Allocation for Content Delivery in the Internet," in submission to *IEEE Trans. on Computers*.

- [10] C.J. Su, L. Tassiulas and V. Tsotras. "Broadcast Scheduling for Information Distribution," *ACM Journal on Wireless Networks*, vol. 5, no. 2, pp 137-147, 1999.
- [11] G. K. Zipf *Human Behaviour and the Principle of Least Effort*, Addison-Wesley Press, Cambridge MA, 1949.
- [12] O. Ercetin, Market-Based Resource Allocation for Content Delivery in the Internet, PhD Dissertation, University of Maryland, College Park, expected completion date 2001.

Personnel supported

- **Faculty:** Leandros Tassiulas
- **Students (PhD) :** Iordanis Koutsopoulos, Ozgur Ercetin.

Publications Supported by Grant #F49620-98-1-0156

- I. Koutsopoulos and L. Tassiulas, "A synchronization-based scheme for simultaneous full- and half-duplex communication in GSM-based MEO mobile satellite networks", *Proceedings of IEEE Globecom*, pp.276-280, 1999.
- I. Koutsopoulos and L. Tassiulas, "Efficient resource utilization through carrier grouping for half-duplex communication in GSM-based MEO mobile satellite networks", *IEEE Transactions on Wireless Communications*, vol.1, no.2, pp.342-352, April 2002.
- I. Koutsopoulos and L. Tassiulas, "A unified framework for handover prediction and resource allocation in non-geostationary mobile satellite networks", *Proceedings IEEE Vehicular Technology Conference Fall*, vol.4, pp.2106-2110, 1999.
- I. Koutsopoulos and L. Tassiulas, "Reliable handover prediction and resource allocation in MEO mobile satellite networks", *Proceedings of IEEE/IMACS Circuits, Systems, Communications Conference*, 1999.
- O. Ercetin and L.Tassiulas, "Push Based Information Delivery in Two Stage Satellite-Terrestrial Wireless Systems," to appear in *IEEE Trans. on Computers*.
- O. Ercetin and L.Tassiulas, "Optimal Competitive Resource Allocation for Content Delivery in the Internet," in submission to *IEEE Trans. on Computers*.
- C.J. Su, L. Tassiulas and V. Tsotras. "Broadcast Scheduling for Information Distribution," *ACM Journal on Wireless Networks*, vol. 5, no. 2, pp 137-147, 1999.
- O. Ercetin, "Market-Based Resource Allocation for Content Delivery in the Internet", PhD Dissertation, University of Maryland, College Park, expected completion date 2001.




Figure 1: Comparison of the cases of single and multiple burst offset values.




Figure 2: Different delay classes within a satellite footprint.

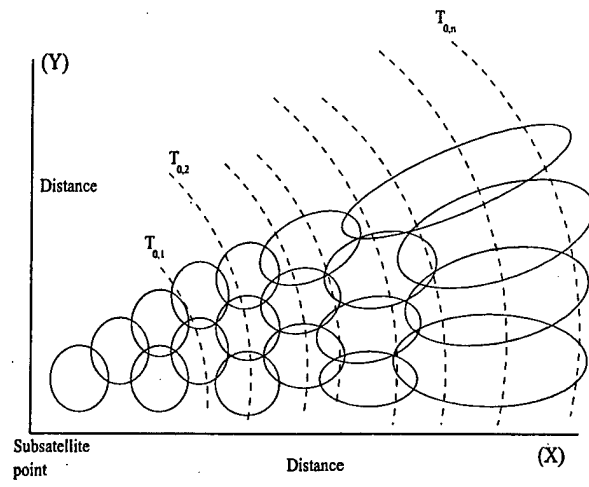


Figure 3: Relative position of delay classes and beams in the satellite footprint.




Figure 4: Illustration of first step of delay class determination algorithm.

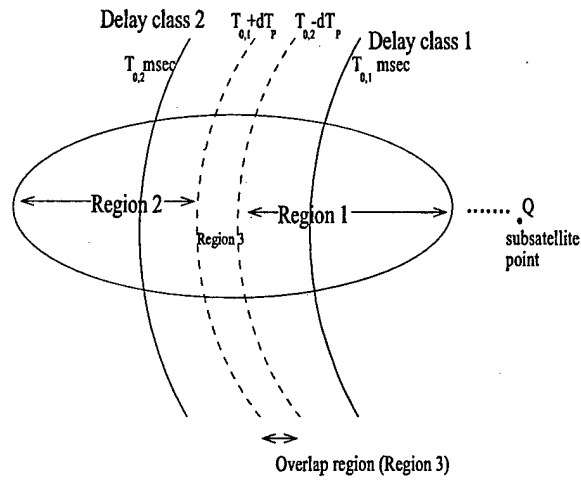


Figure 5: A beam with two delay classes and the corresponding overlap region.

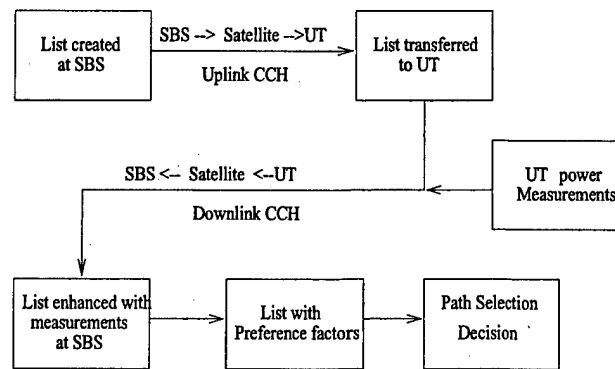


Figure 6: Schematic representation of Path Selection procedure.

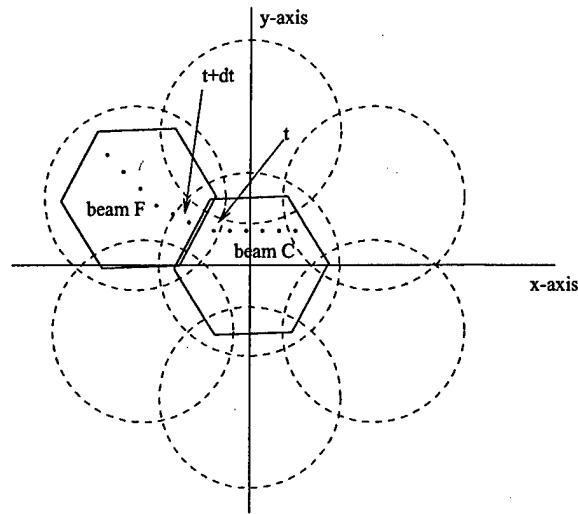


Figure 7: Demonstration of beam handover from the beam located in the satellite nadir (C) to a neighboring beam (F).

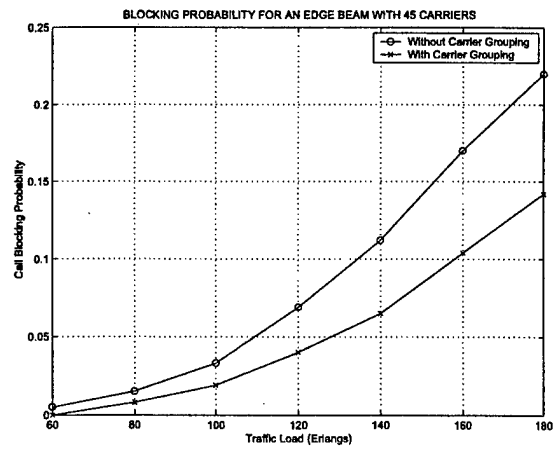
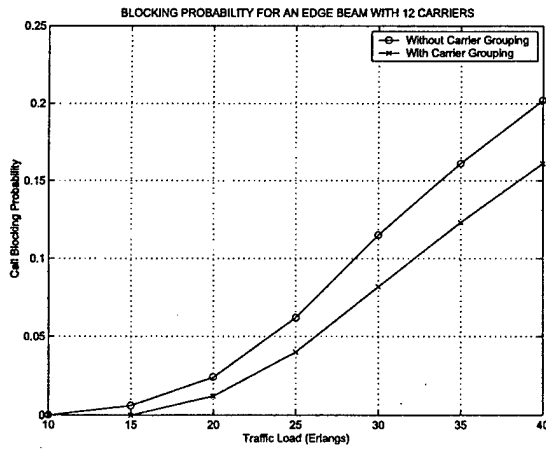



Figure 8: Blocking rates with and without the carrier grouping method, for an edge beam with 12 and 45 carriers.

AGGREGATE SATELLITE, BEAM AND DELAY CLASS HANDOVER RATES DURING BUSY HOUR AT (5,17)



ENTIRE GLOBE : SATELLITE, BEAM AND DELAY CLASS HANDOVER RATES DURING BUSY HOUR


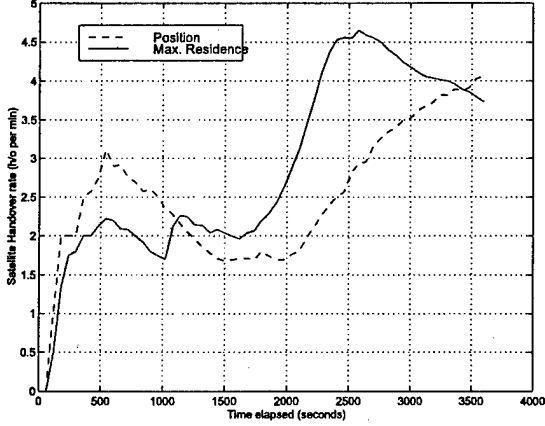


Figure 9: Satellite, beam and delay class handover rates: Handover rates at a square region of 15° longitude and latitude distance from the GW, with $\lambda = 4.87$ calls/sec and handover rates for all calls served by a GW.

SATELLITE HANDOVER RATE WITH POSITION OR MAX. BEAM RESIDENCE CRITERION



BEAM HANDOVER RATE WITH POSITION OR MAX. BEAM RESIDENCE CRITERION




Figure 10: Satellite and beam handover rates under UT position or maximum beam residence time handover criterion.

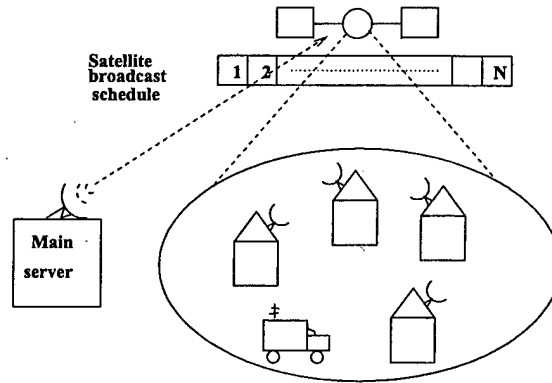


Figure 11: A Single-stage Broadcast Data Delivery System in a Wireless Commment




Figure 12: Multistage Information Distribution

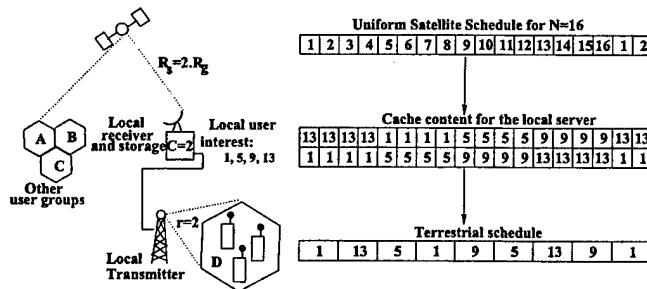


Figure 13: The cache management, and terrestrial schedule at one of the four users of this ground station are equally interested in items 1, 5, 9, 13. The satellite is as fast as the terrestrial channel.




Figure 14: The evolution of the expected user backlog of item i observed by a local station. The upper side of the u -axis refers to the satellite transmissions, while the lower side refers to the terrestrial transmissions.




Figure 15: Illustration for ETT algorithm. The upper side of the u -axis refers to the satellite transmissions while the lower side refers to the terrestrial transmissions. t_1 denotes the first estimated terrestrial transmission slot of item i after k th slot.

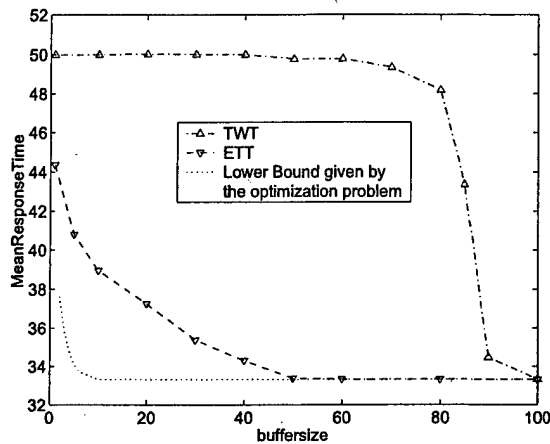


Figure 16: The performances of TWT and ETT. The satellite schedule is uniform with 100 items to broadcast, and the user access probability distribution is Zipf 1 distribution over 100 items broadcast.

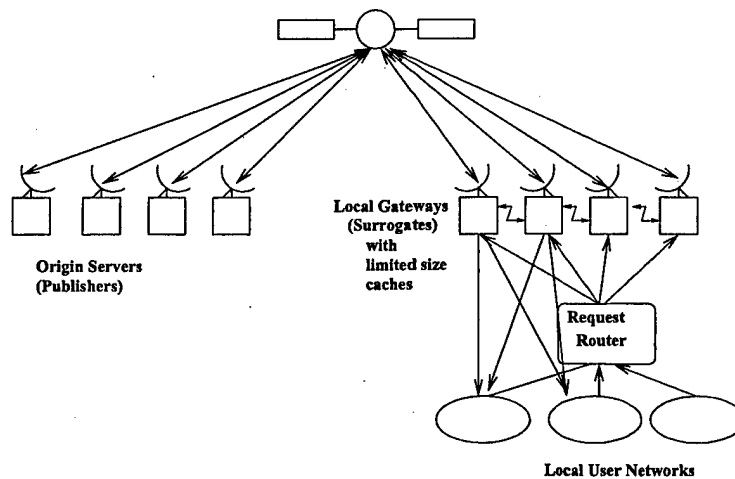


Figure 17: A general multi-stage satellite information retrieval system.

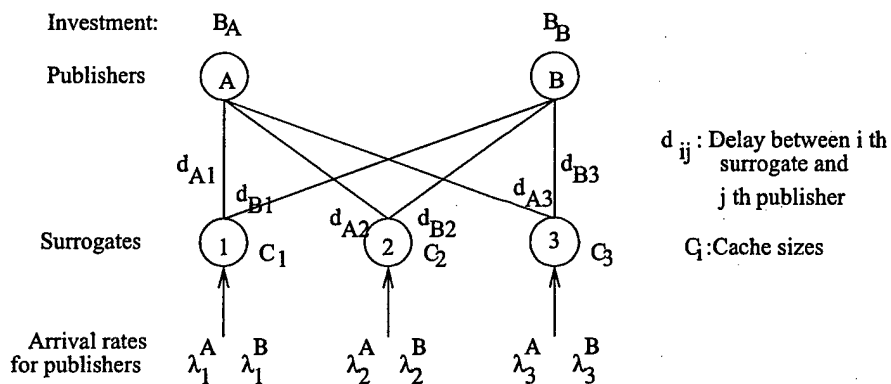


Figure 18: Content delivery system.

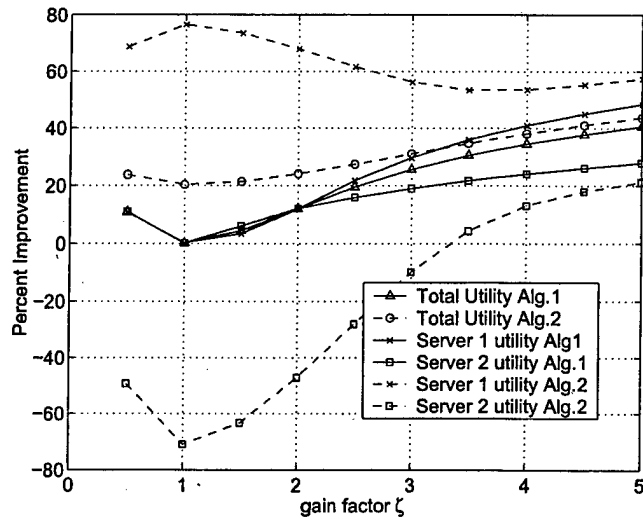


Figure 19: The percentage of improvement of using the game-theoretic algorithm with varying investments over the conventional caching methods. $\alpha_1 = 0.1$, $\alpha_2 = 0.4$. $\zeta_1^1 = \zeta_2^3 = \zeta$. $\zeta_i^j = 1, \forall (i, j) \neq (1, 1), (2, 3)$. $\lambda_1^1 = \lambda_2^1 = \lambda_1^3 = \lambda_2^3 = 3$. $\lambda_i^j = 1, \forall (i, j) \neq (1, 1), (2, 1), (1, 3), (2, 3)$.

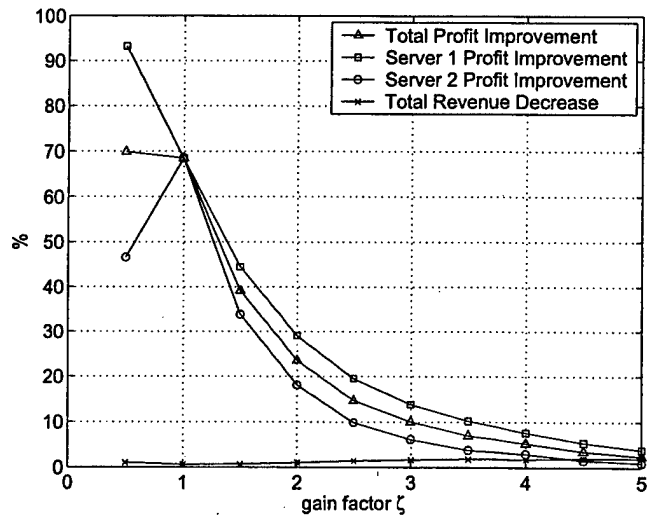


Figure 20: The percentage of improvement of using the game-theoretic algorithm with varying investments over the conventional caching methods. The comparison is made with fixed investments where $B_1 = B_2 = 10$. $\alpha_1 = 0.5$, $\alpha_2 = 0.5$. $\zeta_1^1 = \zeta_2^3 = \zeta$. $\zeta_i^j = 1, \forall (i, j) \neq (1, 1), (2, 3)$. $\lambda_1^1 = \lambda_2^1 = \lambda_1^3 = \lambda_2^3 = 3$. $\lambda_i^j = 1, \forall (i, j) \neq (1, 1), (2, 1), (1, 3), (2, 3)$.