The Concept of Fairness: Definitions and Use in Bandwidth Allocation Applied to Satellite Environment

Igor Bisio, Mario Marchese
University of Genoa

INTRODUCTION

Bandwidth is allocated to entities, such as channels, routers, satellite and wireless stations, servers, computers, and networks, but also to user flows and groups of flows, to reach a specific performance aim, such as minimizing data loss, delay, jitter, and power consumption. Resource allocation (bandwidth, in this case) may be used whenever a common resource must be shared among different entities. In the specific case of satellite environments, which are often affected by noise and fading, bandwidth allocation (better if dynamic) is not only a way to manage channel access and to improve channel utilization but also is traditionally seen as a noise and fading countermeasure, which gives more bandwidth to faded stations and tries to compensate for the penalization introduced by fading while taking bandwidth from the other stations.

Many bandwidth allocation solutions in the literature, independently of the application context, consider the concept of fairness, generically intended as sufficient balance either in the allocated bandwidth among different entities or, often, in the performance offered by the different entities. However, fairness is seldom formally defined. Many scientific articles claim their proposal is fair even if it is in strict contrast with other articles claiming the same. Probably nobody lies and nobody is wrong. They likely use different definitions of fairness.

This article reports on and discusses some definitions of fairness, with the aim of allowing a full comprehension of its meaning and a possible use for the performance evaluation of bandwidth allocation algorithms. The chosen application environment is satellite communication because the authors have practical experience in this field, because in this scenario resource allocation is a delicate issue [1]–[3], and because this environment also allows discussion of the effect of fading. Beyond the shown numerical values, the conclusions will be applicable to other communication environments.

Authors' current address: I. Bisio, M. Marchese, DITEN—Department of Electrical, Electronic and Telecommunications Engineering and Naval Architectures, University of Genoa, Via Opera Pia 13, 16145, Genoa, Italy, E-mails: (igor.bisio@unige.it, mario.marchese@unige.it).
Manuscript received August 12, 2013, and ready for publication September 26, 2013. Review handled by R. Wang.
0885/8985/14/ $26.00

GENERAL CONCEPT OF BANDWIDTH ALLOCATION AND APPLICATION TO THE SATELLITE ENVIRONMENT

BANDWIDTH ALLOCATION

The basic concept of bandwidth allocation is simple, independently of the target of the allocation. An overall amount of bandwidth must be shared among different Z entities. A control mechanism is devoted to this action. The overall bandwidth is C_{tot}. Each entity $z \in [0, Z-1]$ receives a portion C_z of C_{tot}, where $\sum_{z=0}^{Z-1} C_z \le C_{tot}$. Imposing the equality constraint means using all the available capacity. The allocation vector is defined as $\mathbf{C} = (C_0, \ldots, C_z, \ldots, C_{Z-1})$. The control architecture may be supposed to be either centralized, when one entity manages the resources and provides the other ones with a portion of the overall bandwidth, or distributed, when each entity decides its amount of bandwidth on the basis of remote information.

If this generic concept is applied to satellite communications, it takes a special interest, because the bandwidth C_z given to entity Z is not necessarily entirely used to send information. A portion of it may be used to protect information, as should be clear in the following. As a consequence, the net bandwidth for information data is not C_z but rather a lower value. The next session will introduce the concept of real bandwidth to match this bandwidth reduction. The possibility of having a difference between allocated and used bandwidth is typical of satellite environments. This is the main motivation for using satellites as the reference application environment of this article.

SATELLITE ENVIRONMENT

In detail, a generic satellite network is composed of a number (e.g., Z) of earth stations. They are connected through a satellite channel (Figure 1). Each user may request service (e.g., Web page, data transfer, phone call, or audio and video conferencing) by using the satellite channel. To carry out the process, each earth station conveys traffic from the sources and accesses the channel in competition with the other earth

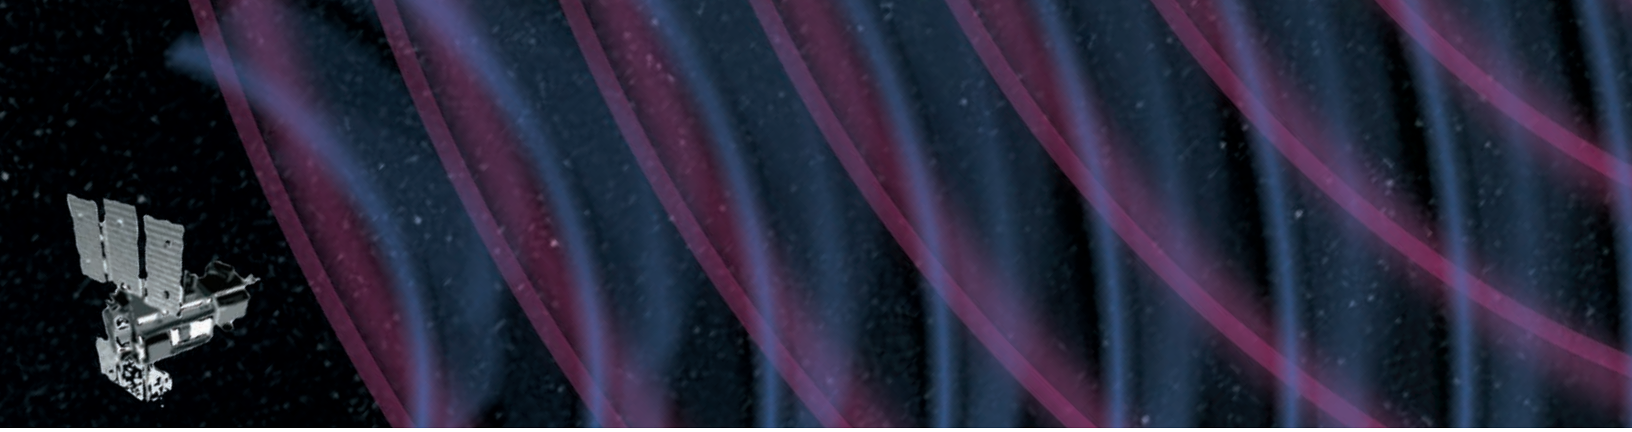

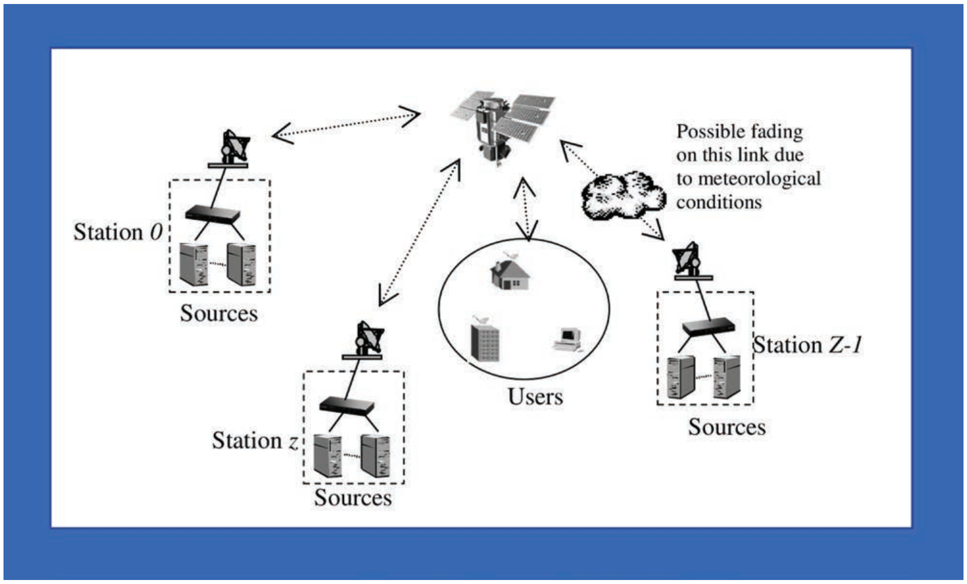

Figure 1.
Satellite network topology.

stations. The earth stations are the entities indicated in the previous discussion. The two terms, station and entity, will be used equivalently in the rest of this article, with a preference for "entity" when its use goes beyond the satellite framework and has a wider application.

Satellite links may be affected by fading and/or noise, as graphically shown for station Z – 1 in Figure 1. Earth stations are modeled as nodes gathering transmission control protocol (TCP)/Internet protocol traffic from the sources, but even if this has a great impact on the shown results, it may be considered an example and it does not affect the generality of the introduced concepts.

ALLOCATED AND USED BANDWIDTH: THE CONCEPT OF BANDWIDTH REDUCTION

Noise and fading corruption typically due to rain are predominant at high frequencies, especially above 10 GHz, and must be compensated to assure efficient services. A widespread compensation solution applied at the physical layer is the employment of forward error correction (FEC) coding schemes, aimed at protecting the integrity of information under different noise and fading conditions by dedicating part of the bits to this aim. Enlarging the number of protection bits means extending the power of FEC code but also reducing the information bit rate. So, the consequence of the use of powerful coding schemes is bandwidth reduction, which may be applied, as in [4], through a proper factor. Numerically, it means that the real bandwidth C_z^{real} available for the zth station is a portion of the nominal bandwidth C_z, which is reduced of a factor β_z, a variable parameter contained in the real numbers interval [0,1]: $C_z^{real} = \beta_z \cdot C_z$. The corresponding vector that contains the bandwidth used for data transmission is $\mathbf{C}^{real} = \left(C_0^{real},\ldots,C_z^{real},\ldots,C_{Z-1}^{real}\right)$.

Even if the idea of having a reduction with respect to the allocated bandwidth, i.e., the bandwidth operatively used by an entity is only a portion of the bandwidth allocated to that entity, has its origin in satellite/wireless communications, it can be applied to other environments. It enlarges the scope of the discussion, because the fairness of an allocation scheme probably should be measured not on the allocated bandwidth but rather on the bandwidth "seen" and used by an entity. The β-based model presented above allows this.

FAIRNESS

GENERAL CONCEPT

The concept of fairness is used in various frameworks, such as economic and social sciences, computer science, and telecommunications, when a limited amount of resources must be simultaneously allocated to different entities.

As said in [5], finding a common definition and a specific index to quantify and compare the degree of fairness of different resource allocation policies so as to avoid ambiguous interpretations is highly desirable. That is particularly important in case of possible reduction of the assigned bandwidth.

Fairness might be considered the same as equality, but this is not correct. Allocating the same amount of bandwidth to each entity may imply a strong performance unbalance among entities. An example concerning satellite networks may help in understanding this distinction. There are two earth stations, and bandwidth is equally shared between them—even if one of them is heavily affected by fading and uses a very small amount of bandwidth for information transport. Bandwidth distribution seems apparently fair but leads to a strong performance unbalance, because the bandwidth actually used by the faded station to transmit information is only a portion of the allocated bandwidth. Is it fair?

Fairness cannot be simply considered as equal resource distribution without taking into account system configurations and conditions and users' expectations (e.g., see [6]). Avoiding ambiguities implies a careful definition of fairness. References [5], [7], and [8] apply the concept of fairness in traditional wired networks and introduce two interesting general-purpose fairness indexes and one fairness concept: the max–min fairness (MMF) index [7], the Jain fairness index [5], and the proportional fairness (PF) concept [8]. They are summarized in the following. In all cases, the real bandwidth ($\beta_z \cdot C_z$) used by a station z, and not the allocated bandwidth C_z, is applied to evaluate fairness and thus consider the impact of bandwidth reduction explicitly. The two values are the same if there is no fading and no attenuation, as typical in wired telecommunications.

MMF INDEX

The concept of MMF was adopted by the Asynchronous Transfer Mode Forum to specify fairness in wired data networks. The concept was proposed and formally defined in [7]. The MMF index I_{MMF}, adapted by using real bandwidth, is defined as follows:

$$I_{MMF} = \frac{\min\left\{\beta_0 \cdot C_0, \ldots, \beta_z \cdot C_z, \ldots, \beta_{Z-1} \cdot C_{Z-1}\right\}}{\max\left\{\beta_0 \cdot C_0, \ldots, \beta_z \cdot C_z, \ldots, \beta_{Z-1} \cdot C_{Z-1}\right\}} \quad (1)$$

The denominator and nominator of (1) represent the maximum and minimum bandwidth, respectively, provided to the involved entities. The I_{MMF} value is bounded and ranges between 0 and 1. The latter is the ideal value, because it implies that all entities can use the same amount of bandwidth (the same $\beta_z \cdot C_z$, $\forall i \in \left[0, Z-1\right]$, in this article), and this is the real aim of max–min-based allocations. Actually, max–min-based allocations look for the allocation that approaches to 1 the MMF index, so making close minimum and maximum real bandwidth allocations $\beta_z \cdot C_z$, $z \in \left[0, Z-1\right]$. I_{MMF} is independent of scale and unit of measurement. Intuitively, any change of bandwidth allocation should imply a variation of the fairness index, but bandwidth changes do not necessarily imply a variation of the MMF index. This is a limitation of I_{MMF}. For example, if two different real bandwidth allocations provide the same minimum and maximum capacity assignations, the index is the same. In short, the variation of I_{MMF} is governed only by the ratio between minimum and maximum real bandwidth allocations, not by each product $\beta_z \cdot C_z$, $z \in \left[0, Z-1\right]$. Even if the MMF index is a widely adopted fairness index, as described in [6], it has a clear drawback because it neglects the distribution of the real bandwidth among the entities.

JAIN FAIRNESS INDEX

An alternative fairness definition comes from Jain fairness index [5] I_{JF}. Considering the bandwidth assigned to the entities, the Jain fairness index is defined as

$$I_{JF} = \frac{\left|\sum_{z=0}^{Z-1} \beta_z \cdot C_z\right|^2}{Z \cdot \sum_{z=0}^{Z-1} \left(\beta_z \cdot C_z\right)^2} \quad (2)$$

where Z is the overall number of entities. It has the following properties [5]: It is independent of the scale and of the unit of measurement, it is bounded, and its value ranges between $\frac{1}{Z}$ and 1 (0.5 and 1 if $Z = 2$). In a completely fair system, the fairness index is equal to 1; it is equal to $\frac{1}{Z}$ if the system is completely unfair. When a real bandwidth allocation changes slightly, the Jain fairness index also varies. I_{JF} measures the degree of distribution of the overall real bandwidth $\sum_{z=0}^{Z-1} \beta_z \cdot C_z$ among the different Z entities.

COMMENTS

A simple numerical example may help understand the real meaning of the two mentioned indexes and the role of the fading factor β_z. There are two entities (earth stations, in this case), identified through the indexes 0 and 1. The available overall bandwidth is 4 Mbps. The values of the MMF and Jain fairness indexes are reported in Figure 2 versus the bandwidth C_0 available for station 0, both not affected by fading and by varying the fading factor β_1 of station 1. Imposing the equality in the constraint $\sum_{z=0}^{Z-1} C_z \le C_{tot}$, $C_1 = 4 - C_0$ Mbps. The aim of Figure 2 is to check the index behavior by changing the bandwidth balance between the two earth stations for different fading levels.

Both fairness indexes are polarized by the fading factor β_1. When $\beta_1 = 1$, both indexes have the maximum value 1 when the real bandwidth distribution is balanced ($C_0 = C_1 = 2$ Mbps). If β_1 decreases, the maximum value of both indexes moves to configurations that privilege the faded station. In numbers, the Jain fairness index I_{JF} is 1 when [C_0 = 1.71, C_1 = 2.29] if β_1 = 0.75, when [C_0 = 1.33, C_1 = 2.67] if β_1 = 0.5, and when [C_0 = 0.8, C_1 = 3.2] if β_1 = 0.25. Similar numbers may be obtained from Figure 2 for I_{MMF}. It is interesting to highlight the differ-

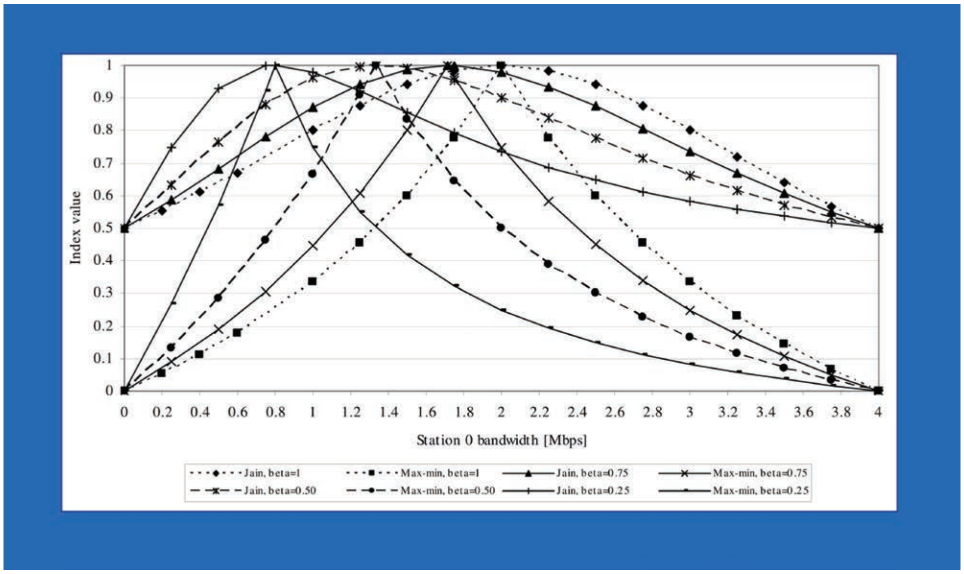

Figure 2.
MMF and Jain fairness indexes versus station 0 bandwidth by varying the fading factor.

ent shape of the curves produced by the two indexes. Independently of the different value ranges, which are $\left[0.5,1\right]$ for I_{JF} and $\left[0,1\right]$ for I_{MMF}, the Jain fairness index decreases less quickly than the MMF index even if, fixing the bandwidth of station 0, the two indexes are in agreement concerning the order when β_1 varies, i.e., $I_{JF}\big|_{\beta_1=1} > I_{JF}\big|_{\beta_1=0.75} > I_{JF}\big|_{\beta_1=0.5} > I_{JF}\big|_{\beta_1=0.25}$ and $I_{MMF}\big|_{\beta_1=1} > I_{MMF}\big|_{\beta_1=0.75} > I_{MMF}\big|_{\beta_1=0.5} > I_{MMF}\big|_{\beta_1=0.25}$. The different shapes of the curves provided by the two indexes may be seen into two ways: both fixing β_1 and fixing bandwidth. Let us fix β_1: A slight unbalance of bandwidth assignation implies a slight decrease of the Jain fairness index but a drastic decrease of the MMF index. For example, fixing β_1 to 1, the behavior is the same by using other values for β_1. When [C_0 = 2.25, C_1 = 1.75], I_{JF} is 0.98 but I_{MMF} goes down to 0.78; when [C_0 = 2.5, C_1 = 1.5], I_{JF} is 0.94 but I_{MMF} decreases to 0.6; and when [C_0 = 2.75, C_1 = 1.25], the Jain fairness index is still close to 1 (0.88) but the MMF index has more than halved its original value and is 0.45.

If the bandwidth is fixed (e.g., [C_0 = C_1 = 2)]), a change of β_1 does not affect the two indexes in the same way. If β_1 = 0.75, the Jain fairness index is still 0.98 while the MMF index decreases to 0.75; if β_1 = 0.5, the Jain fairness index is 0.9 but the MMF index is down to 0.5. The same trend is valid for β_1 = 0.25.

UTILITY-BASED FAIRNESS

A particular bandwidth allocation may be considered fair or unfair not only on the basis of the amount of provided bandwidth (C_z) and of the channel status (β_z), as done for I_{JF} and I_{MMF}, but also dependently on the use of the bandwidth and of the possible user/network revenue. A proper index in this case is associated with the notion of utility-based fairness (UBF), which has been introduced in [8] and has its main strength in the flexibility of the concept that can be customized for a variety of different applications [8]. UBF can be defined by introducing a utility value for each considered entity as a function of the allocated bandwidth and of other parameters. It is aimed at measuring the expectations of each entity in terms of revenue, quality, and prize. Utility functions are supposed to be concave and denoted as $U_z\left(C_z,\beta_z,\cdot\right)$. With Z as the number of entities, a UBF allocation $\mathbf{C}^* = \left(C_0^*,\ldots,C_z^*,\ldots,C_{Z-1}^*\right)$ is defined as the allocation that maximizes the sum of the utility functions over the overall number of entities $\sum_{z=0}^{Z-1} U_z\left(C_z,\beta_z,\cdot\right)$. Having an index that quantifies the "distance" of a given allocation from the UBF allocation may be useful. The aim is to measure how much the utility of the overall system is close to its maximum. Fixed $U_z\left(\cdot\right), \forall z$, the UBF index I_{UBF}, originally introduced in this article, is defined as in (3). It may assume values in the interval $\left[0,1\right]$, where 1 corresponds to the UBF.

$$I_{UBF} = \frac{\sum_{z=0}^{Z-1} U_z\left(C_z,\beta_z,\cdot\right)}{\sum_{z=0}^{Z-1} U_z\left(C_z^*,\beta_z,\cdot\right)} \quad (3)$$

I_{UBF} represents a family of fairness indexes, each of them defined by one specific utility function $U_z\left(\cdot\right)$.

PROPORTIONAL FAIRNESS

An operative example of UBF is PF, proposed in [8]. The utility function in PF is the logarithm of another function, $U_z\left(C_z,\beta_z,\cdot\right) = \ln\left(f_z\left(C_z,\beta_z,\cdot\right)\right)$, $f_z\left(C_z,\beta_z,\cdot\right) > 0$, so the quantity to maximize to get a proportional fair allocation is the product $\prod_{z=0}^{Z-1} f_z\left(C_z,\beta_z,\cdot\right)$, with $\ln\left(\cdot\right)$ as a strictly increasing function. Directly from I_{UBF}, the PF index I_{PF} may be defined as follows:

$$I_{PF} = \frac{\sum_{z=0}^{Z-1} \ln\left(f_z\left(C_z,\beta_z,\cdot\right)\right)}{\sum_{z=0}^{Z-1} \ln\left(f_z\left(C_z^*,\beta_z,\cdot\right)\right)} = \frac{\ln\left(\prod_{z=0}^{Z-1} f_z\left(C_z,\beta_z,\cdot\right)\right)}{\ln\left(\prod_{z=0}^{Z-1} f_z\left(C_z^*,\beta_z,\cdot\right)\right)} \quad (4)$$

BANDWIDTH ALLOCATION SOLUTIONS

PHYSICAL CONSTRAINT

The aim is to briefly summarize some bandwidth allocation solutions appearing in the literature so to allow computing of the indexes defined in the previous section and performance of a comparison. The equality constraint $\sum_{z=0}^{Z-1} C_z = C_{tot}$ is imposed for all the methods as typically done in the literature, where the performance metrics improve when the allocated bandwidth increases. This happens also in this article, with $P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)$ as the used reference metric for the listed allocation methods, as described below. $P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)$ is the loss probability of the TCP packets, and its analytical expression is defined in [9]. It is an increasing function of the variable C_z. Coherently, the utility function $U_z\left(C_z, \beta_z, \cdot\right)$ in (3) is set to $\ln\left(\frac{1}{P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)}\right)$ in the tests. Obviously, $f_z\left(C_z, \beta_z, \cdot\right)$ in (4) is $\frac{1}{P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)}$, which is a concave function, as required in [8] and where it has been shown that if $f_z\left(C_z, \beta_z, \cdot\right)$ is concave, $\ln\left(f_z\left(C_z, \beta_z, \cdot\right)\right)$ also is concave.

FIXED ALLOCATION

With the fixed bandwidth allocation method (FIX), the bandwidth allocator assigns the same capacity to each station independently of noise, fading and traffic conditions: $C_z = \frac{C_{tot}}{Z}$.

HEURISTIC ALLOCATION

Assuming both the fading factor β_z and the traffic load, expressed as the number of active connections N_z offered at earth station $z \in \left[0, Z-1\right]$, are known, the bandwidth provided to the zth station is computed as a weighted portion of the overall available bandwidth $C_z = k_z \cdot C_{tot}$ using heuristic allocation (HEU). The weight k_z is set to $N_z / \beta_z \left(\sum_{j=0}^{Z-1} N_j / \beta_j\right)^{-1}$.

The bandwidth assigned to a station increases with the traffic offered to the station and decreases with the bandwidth reduction.

VALUE FUNCTION

The value function (VALUE) bandwidth allocation strategy [10] distributes the bandwidth by minimizing the sum of the single functions $P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)$. In short, the bandwidth is allocated by minimizing the function $\sum_{z=0}^{Z-1} P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)$.

NASH BARGAIN SOLUTION

The Nash bargain solution (NBS), deeply investigated in [8], is based on the Nash bargaining problem, which originated from the bargaining theory [11]. NBS maximizes the "social benefit," which is the product of utility functions, chosen here as stated in the physical constraints discussion of the bandwidth allocation solution earlier in this article. In consequence, the bandwidth is allocated by maximizing the function $\prod_{z=0}^{Z-1} \frac{1}{P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)}$. Choosing the same utility functions to compute the UBF index and to allocate bandwidth for NBS is not mandatory. The choice seems reasonable and coherent with the performance metric used for the other allocation schemes shown in this article, but any alternative of concave functions would be acceptable.

UTOPIA MINIMUM DISTANCE

Utopia minimum distance (UMD) has been thought out for a fully competitive environment, and it is aimed at approaching ideal performance, which theoretically happens when each station has the full availability of all the channel bandwidth. UMD minimizes the square of the Euclidean distance between the performance vector $\left\{P_0^{TCP-loss}\left(C_0, \beta_0, \cdot\right), \ldots, P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right), \ldots, P_{Z-1}^{TCP-loss}\left(C_{Z-1}, \beta_{Z-1}, \cdot\right)\right\}$ and the ideal, not feasible, performance vector obtained by setting $C_z = C_{tot},\ \forall z \in \left[0, Z-1\right]$. The bandwidth is allocated by minimizing the function $\sum_{z=0}^{Z-1}\left(P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right) - P_z^{TCP-loss}\left(C_{tot}, \beta_z, \cdot\right)\right)^2$.

INDEXES COMPUTATION AND COMPARISON

NUMERICAL COMPARISON

The bandwidth allocation methods FIX, HEU, VALUE, NBS, and UMD are used to compute the fairness indexes introduced previously. The comparison allows better understanding of the meaning of the different indexes and a check of the effect of specific resource allocation criteria on fairness. Only TCP traffic is considered, through the metric $P_z^{TCP-loss}\left(C_z, \beta_z, \cdot\right)$ referenced previously. The following bandwidth reduction levels taken from [4] are used: $\beta \in \left(0.15625,\ 0.3125,\ 0.625,\ 0.8333,\ 1\right)$. The overall bandwidth available C_{tot} is set to 4 Mbps, and the TCP buffer size is set to 10 packets of 1500 bytes for each earth station. The TCP round-trip time (RTT) is considered fixed and equal to 100 ms for all the stations. It can represent a medium earth orbit satellite, but it is just an example: RTT numerical values have no impact on the behavioral trends of the fairness indexes. Two earth stations ($Z = 2$) have been taken into account (station 0 and station 1), and the number of active TCP sources is set to $N_z = 10$, $z = \{0,1\}$. Station 0 is supposed to always be in clear sky ($\beta_0 = 1$), while station 1 varies its condition in the tests: $0 \le \beta_1 \le 1$, $\beta_1 \in \left(0.15625,\ 0.3125,\ 0.625,\ 0.8333,\ 1\right)$. The number of earth stations is limited to two to simplify the interpretation of the fairness index value.

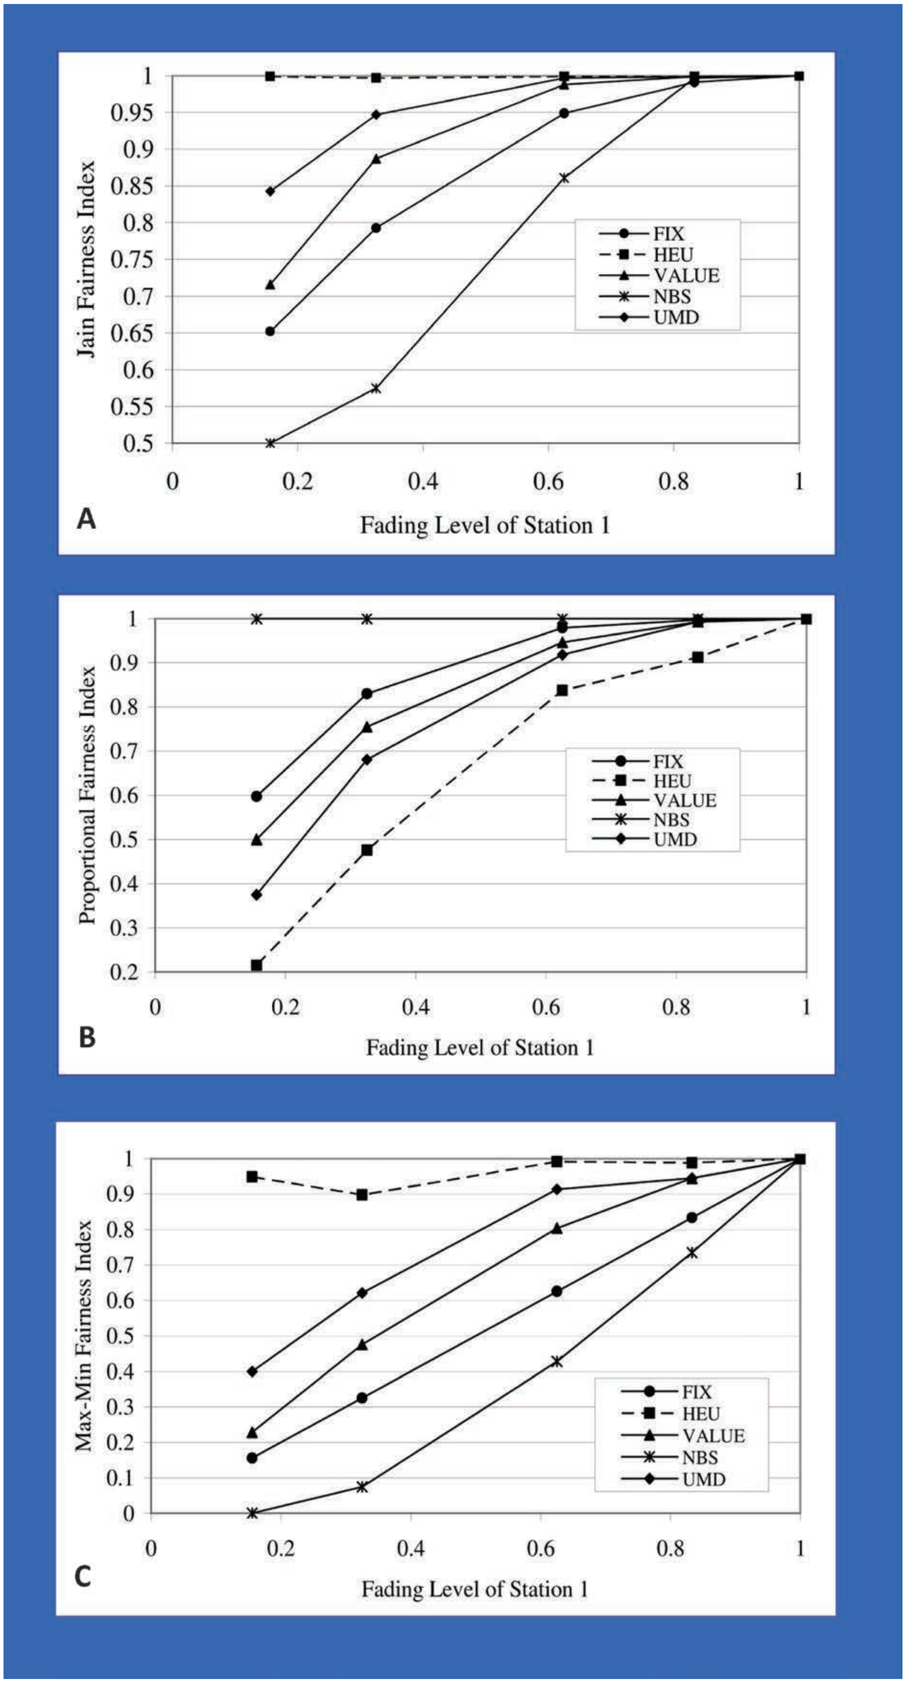

Figure 3.
(A) Jain fairness index. (B) PF index. (C) MMF index.

FAIRNESS INDEXES COMPUTATION

Figure 3a shows the values of the Jain fairness index versus the fading level of station 1. The HEU method is the fairest. Its Jain fairness index is almost constant independently of the bandwidth reduction level. All the other methods have an increasing behavior with β_1. UMD allows high Jain fairness index values, as does VALUE. NBS is the least fair from the Jain fairness index viewpoint. Actually, even if neither VALUE nor UMD uses the Jain fairness index as a guideline for the design, their cost functions automatically reach a sort of fading-level–weighted bandwidth distribution, which increases the Jain fairness index. However, the NBS cost function has a different structure that, due to the utility function choice performed in this article, operatively is the same as maximizing the PF index. This is clear in Figure 3b, where the PF index is shown versus the fading level of station 1. The MMF index is shown in Figure 3c. The same general comments reported for Figure 3a are valid. The numerical effect of the MMF and Jain fairness indexes has been already compared in Figure 2.

CONCLUSIONS

Starting from the literature in the field, the article has discussed and formalized the concept of fairness, as well as reporting on specific evaluation indexes. Fairness indexes have been computed for some resource allocation methods applied to satellite systems taken in the literature, giving numerical examples on which to comment. The analysis has shown that the fairness concept has different meanings and, consequently, must be measured through different indexes. Allocation algorithms may be fair concerning one index and unfair concerning another one (the case of NBS is evident from this viewpoint). This aspect underlines the importance of having a precise definition of fairness within a scientific work to avoid ambiguities. In summary, the MMF and

Jain fairness indexes' results are in agreement concerning the order when the fading factor varies and the bandwidth is fixed (meaning given bandwidth C, if $\beta' > \beta''$, $I_{JF}\big|_{\beta',C} > I_{JF}\big|_{\beta'',C}$ and $I_{MMF}\big|_{\beta',C} > I_{MMF}\big|_{\beta'',C}$), but the slopes of the curves provided by the two indexes are different, also from the concavity viewpoint, and I_{JF} increases and decreases less quickly by varying bandwidth. Moreover, the MMF index neglects the bandwidth distribution among stations, while the Jain fairness index measures the degree of bandwidth distribution. The PF index has a different aim: it not only considers the amount of provided bandwidth and the channel status concerning fading but also weights the bandwidth use and the possible user revenue through a utility function. Because the objectives are different, the results provided by allocation schemes that take PF as a reference (also implicitly) provide different results from the algorithms that have inspiration from MMF and/or Jain fairness.

It must be underlined that the index evaluation is performed after the fact in this article; i.e., the index is computed after bandwidth allocation and has no role in the allocation. Nevertheless, the mentioned concepts of fairness may be used as guidelines for the design of bandwidth allocation strategies including preallocation and/or real-time allocation. This can be done by employing the fairness index definitions as functions to be optimized. The introduced indexes may be also a valid operative help in this direction. Fairness is not the only metric to assess the performance of a bandwidth allocation scheme. Efficiency, which also needs clear definition, is an important metric, for instance. Links and possible tradeoffs between fairness and efficiency need adequate study and insights and will be the object of future research.◆

REFERENCES

[1] Nishiyama, H., Tada, Y., Kato, N., Yoshimura, N., Toyoshima, M., and Kadowaki, N. Toward optimized traffic distribution for efficient network capacity utilization in two-layered satellite networks. *IEEE Transactions on Vehicular Technology,* Vol. 62, 3 (Mar. 2013), 1303–1313.

[2] Bisio, I., and Marchese, M. Power saving bandwidth allocation over GEO satellite networks. *IEEE Communications Letters,* Vol. 16, 5 (May 2012), 596–599.

[3] Nishiyama, H., Kudoh, D., Kato, N., and Kadowaki, N. Load balancing and QoS provisioning based on congestion prediction for GEO/LEO hybrid satellite networks. *Proceedings of the IEEE,* Vol. 99, 11 (Nov. 2011), 1998–2007.

[4] Celandroni, N., Davoli, F., and Ferro, E. Static and dynamic resource allocation in a multiservice satellite network with fading. *International Journal of Satellite Communications,* Vol. 21, 4–5 (July–Oct. 2003), 469–488.

[5] Jain, R., Chiu, D., and Hawe, W. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. DEC Res. Rep. TR-301, Sept. 1984.

[6] Dianati, M., Shen, X., and Naik, S. A new fairness index for radio resource allocation in wireless networks. In *Proceedings of the IEEE Wireless Communications and Networking Conference* (WCNC), Mar. 2005, 712–717.

[7] Bertsekas, D., and Gallager, R. *Data Networks.* Englewood Cliffs, NJ: Prentice Hall, 1992.

[8] Yaïche, H., Mazumdar, R. R., and Rosemberg, C. A game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Transactions on Networking,* Vol. 8, 5 (Oct. 2000), 667–678.

[9] Bisio, I., and Marchese, M. Analytical expression and performance evaluation of TCP packet loss probability over geostationary satellite. *IEEE Communications Letters,* Vol. 8, 4 (Apr. 2004), 232–234.

[10] Bolla, R., Davolim F., and Marchese, M. Adaptive bandwidth allocation methods in the satellite environment. In *Proceedings of the Conference on Computer Communications* (ICC), Helsinki, Finland, June 2001, 3183–3190.

[11] Muthoo, A. *Bargaining Theory with Applications.* Cambridge, United Kingdom: Cambridge University Press, 1999.