

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
18 February 2010 (18.02.2010)

PCT

(10) International Publication Number
WO 2010/019613 A1

(51) International Patent Classification:
H04W 16/28 (2009.01)

(21) International Application Number:
PCT/US2009/053466

(22) International Filing Date:
11 August 2009 (11.08.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/088,712 13 August 2008 (13.08.2008) US
12/538,733 10 August 2009 (10.08.2009) US

(71) Applicant (for all designated States except US): **NTT DOCOMO, INC.** [JP/JP]; Sanno Park Tower, 11-1, Nagata-cho, 2-chome, Chiyoda-ku (JP).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **CAIRE, Giuseppe** [IT/US]; 1948 Huntington Drive, South Pasadena, CA 91030 (US). **RAMPRASHAD, Sean, A.** [GB/US]; 960 Terrace Drive, Los Altos, CA 94024 (US). **PAPADOPOULOS, Haralabos** [GR/US]; 390 South Cypress Avenue, San Jose, CA 95117 (US).

(74) Agents: **VINCENT, Lester, J.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: A METHOD OF COMBINED USER AND COORDINATION PATTERN SCHEDULING OVER VARYING ANTENNA AND BASE-STATION COORDINATION PATTERNS IN A MULTI-CELL ENVIRONMENT

(57) Abstract: A method and system is disclosed herein for scheduling in a multi-cell environment. In one embodiment, the system comprises a plurality of user terminals; a plurality of base stations to communicate with the plurality of user terminals, wherein the plurality of antennas from the base stations are grouped into coordination patterns that specify which antennas operate together as a cluster to coordinate transmissions with a subset of the plurality of user terminals, and further wherein defined pairs of the antenna sets and associated user sets vary over transmission resources based on coordination pattern changes such that different sets of antennas in the plurality of antennas are grouped together on different resources to coordinate transmissions to different subsets of users; and at least one controller to execute at least one scheduling algorithm to schedule transmissions with user terminals by specifying which user terminal to schedule for service by which sets of antennas and on which transmission resource and with what transmission rate.



WO 2010/019613 A1

A METHOD OF COMBINED USER AND COORDINATION PATTERN SCHEDULING OVER VARYING ANTENNA AND BASE-STATION COORDINATION PATTERNS IN A MULTI-CELL ENVIRONMENT

PRIORITY

[0001] The present patent application claims priority to and incorporates by reference the corresponding provisional patent application serial no. 61/088,712, titled, “A Method of Combined User and Coordination Pattern Scheduling over Varying Antenna and Base-Station Coordination Patterns in a Multi-Cell Environment,” filed on August 13, 2008.

RELATED APPLICATIONS

[0002] The present application is related to the following applications: United States Patent Application No. 12/538,729, filed August 10, 2009, titled “A Variable Coordination Pattern Approach for Improving Performance in Multi-Cell or Multi-Antenna Environments” and United States Patent Application No. 12/538,739, filed August 10, 2009, titled “A Method for Varying Transmit Power Patterns in a Multi-Cell Environment”, concurrently filed herewith.

FIELD OF THE INVENTION

[0003] Embodiments of the present invention relates to the field of wireless transmission in a multi-cell wireless environment; more particularly, embodiments of the present invention relate to scheduling in a wireless communication system by varying user terminal rates based on some criteria (e.g., fairness) while varying coordination patterns.

BACKGROUND OF THE INVENTION

Operation of Cells and Multi-Cells and Cell Clusters

[0004] For purposes herein, a “cell” is used to describe a geographic area serviced by a set of (number of) transmit antennas. Such antennas can reside on one or more base-stations. They may be co-located at one point, or geographically spread over the cell. They may not even physically reside in the geographic

boundary of the cell, though it is often the case they are physically located in the cell. For downlink traffic, this set of antennas services user terminals in the cell by jointly transmitting signals to such users. The transmitted signals are produced and controlled by a single common physical layer mechanism. That is, within a given “cell” the transmit antennas are coordinated. Such a common physical layer mechanism can be implemented by a number of processing entities that may be spread over a number of base stations, but which are jointly controlled to achieve a common result. The mechanism may also be implemented by a single processing entity. A classic example of such “cells” is shown in Figure 1. Here the antennas are collocated at a base-station. Referring to Figure 1, user terminals within classic hexagonally shaped cells are mapped to the base-station site that is “geographically closest” (e.g., the center of the cell). This closest mapping results in the classic hexagonal pattern. For example, Cell 1 consists of a central set of 4 antennas at a single base-station, BS1, supporting a group of users including but not limited to user1, user2, user3, and user4. This “geographically closest” station rule makes sense in a model in which the received signal energy a user gets from any base-station (or antenna) decreases monotonically with the distance from that station, and by the same mathematical function for every station. In general, with shadowing and other effects, the cell boundaries will not conform to such a regular structure. In this classic multi-cell scenario, which has a fixed classic cell structure mapping users on a one-to-one basis with base-stations, a single base-station transmits only to users assigned to its cell.

[0005] It is well known that if neighboring base-stations use the same transmission resource (e.g., send signals on the same frequency band at the same time) that users in a cell will experience interference from other cells. Such “inter-cell interference” (ICI) can be quite extreme near the edges of cells, thus limiting performance for users in such areas. This is a classic problem with any cell structure. For example, user4 and user5 in Figure 1 will experience high interference levels relative to their received signal levels. In contrast, users such as user1 and user2 may see less interference and less effect from a multi-cell environment. Nonetheless, all users do experience ICI.

[0006] A scheduling algorithm, such as a proportionally fair scheduling algorithm, can be used to ensure that users such as user1 and user2 are balanced fairly (or a desired fashion) with users seeing less favorable channel conditions such as user4 and user5. Without such a mechanism, if the system simply wanted to deliver the maximum per cell throughput, it would simply transmit only to users with the best channel conditions, severely penalizing other users such as user4 and user5.

[0007] To alleviate the detrimental effects of ICI, one can assume that cells coordinate their transmissions. One such way of doing so is to make sure neighboring base-stations do not use the same frequency at the same time. This results in the classic frequency reuse approach of a cellular system. Another way is to allow neighboring cells to use the same frequency at the same time, but to select signals for transmission that result in low ICI. One way to do so is to have a single underlying physical layer mechanism jointly control antennas from neighboring base-stations. In the extreme, one mechanism controls all transmit antennas over all stations, and the system reduces essentially to a single distributed antenna system with a common MIMO (downlink) broadcast (shared) channel. Here a user such as user4 in cell 1 is in fact serviced by signals coming from all antennas. In such an approach, ICI can be significantly reduced, even set to zero by Multi-user MIMO techniques such as Linear Zero-Forced Beamforming. In fact, in such an extreme the boundaries of “cells” have no meaning, and there are no longer uncoordinated signals crossing cell boundaries creating ICI. However, such a system is impractical.

[0008] A partial coordination may be more practical as shown in Figure 2. Referring to Figure 2, groups of cells (in this case three neighboring cells) can coordinate with each other. As illustrated, cells 1, 2 and 3 coordinate transmissions, cells 4, 5 and 6 coordinate transmissions, cells 7 and 8 coordinate, and cell 10 operates without coordinating with others. Each of these groups of three cells act as a single “cell”, or what is termed a “cluster”. If such coordination between base-stations is done correctly, users such as user4 and user5 can be made to see both less interference, and possibly better signal terms because they are no longer at a boundary with a neighboring cluster. They have now the option to use

transmissions from up to 12 antennas over 3 remote locations. Thus, the system becomes more efficient. In addition, the scheduling of users in the cluster of cells 1, 2 and 3 can be done together.

[0009] However, because coordination is partial, such a system will inherently always have boundaries where users that see less favorable conditions. This is the case of users user6 and user7 in Figure 2 which are at the boundary of two coordinating clusters.

SUMMARY OF THE INVENTION

[0010] A method and system is disclosed herein for scheduling in a multi-cell environment. In one embodiment, the system comprises a plurality of user terminals; a plurality of base stations to communicate with the plurality of user terminals, wherein the plurality of antennas from the base stations are grouped into coordination patterns that specify which antennas operate together as a cluster to coordinate transmissions with a subset of the plurality of user terminals, and further wherein the defined pairs of the antenna set and its associated user set vary over transmission resources based on coordination pattern changes such that different sets of antennas in the plurality of antennas are grouped together on different resources to coordinate transmissions to different subsets of users; and at least one controller to execute at least one scheduling algorithm to schedule transmissions with user terminals by specifying which user terminal to schedule for service by which sets of antennas and on which transmission resource and with what transmission rate.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

Figure 1 illustrates user terminals within classic hexagonal cells mapped to a base-station in center of the cell.

Figure 2 is a diagram of a wireless communication system configured as a multi-cell environment with coordination among groups of cells.

Figure 3 illustrates a multi-cell case with coordination among groups of cells with a different coordination pattern than is shown in Figure 2.

Figure 4 illustrates constituent schedulers that may be used in the coordination pattern of Figure 2.

Figure 5 illustrates constituent schedulers that may be used in the coordination pattern of Figure 3.

Figure 6 illustrates information is shared between constituent schedulers in Figure 4 and Figure 5 where some scheduler serves the same (common) users over different coordination patterns (such patterns exist on different transmission resources).

Figure 7 illustrates a communication system where some of the information necessary for each constituent scheduler to perform its own scheduling is stored centrally. .

Figure 8 is a flow diagram of one embodiment of a scheduling process which uses a central repository.

Figure 9 is a diagram of one embodiment of a communication system that includes a coordination controller that decides what coordination pattern to use.

Figure 10 is a flow diagram of another embodiment of a scheduling process.

Figure 11 is a block diagram of a communication system that uses scheduling entities that handle (operate over) multiple clusters.

Figure 12 is a block diagram of one embodiment of a controller as described herein.

DETAILED DESCRIPTION OF THE PRESENT INVENTION

[0012] A method and communication system for operating a user scheduling algorithm within a multi-cell system are described. In the multi-cell systems of interest, different subsets of antennas are used to transmit to different subsets of

users, antennas within each subset can be coordinated in their transmissions, and the subset definitions can change over different transmission resources. A subset of antennas and the associated subset of users served by such antennas constitute what is termed an antenna-user pair. The totality of such pairs for a given transmission resource constitutes a coordination pattern. Such a pattern exists on a transmission resource, e.g. at a given time, and/or at a given frequency, and/or given tones, and/or different codes in a Code Division Multiple Access (CDMA) system. In all embodiments, the system varies the coordination patterns with transmission resources, such that at least two such patterns exist. For example, the coordination patterns can change over time (e.g., periodically) and/or over frequency (e.g., only some patterns on some frequencies) or some combination of both, or over any units of transmission resources, as described in United States Patent Application No. 12/538,729, filed August 10, 2009, titled "A Variable Coordination Pattern Approach for Improving Performance in Multi-Cell or Multi-Antenna Environments", which is a non-provisional application of Application No. 61/088,714, filed on Aug 13, 2008. Such changes in coordination patterns may occur in a known, and/or pre-determined controlled, fashion over time and/or over frequency. The operation of the user terminal scheduling algorithm described herein is for use with such a system whereby it takes particular advantage of such changes in coordination patterns. Its novel implementation in various embodiments is also tied directly to the physical embodiment of a system with varying coordination patterns. For example, the scheduling algorithm may be divided into a number of constituent scheduling entities which exchange sufficient information to make equivalent scheduling decisions to what a single global entity (knowing all information) would make. In one embodiment, the scheduling algorithm also directs (decides) what coordination pattern to use on what resource in an adaptive fashion.

[0013] Techniques disclosed herein include joint, or semi-joint, scheduling algorithms used to operate with a communication system that varies the multi-cell coordination topology. The scheduling algorithm itself can also influence the way the multi-cell system coordination topology is varied.

[0014] In the following description, numerous details are set forth to provide a more thorough explanation of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

[0015] Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0016] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0017] The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or

reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0018] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0019] A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory (“ROM”); random access memory (“RAM”); magnetic disk storage media; optical storage media; flash memory devices; etc.

Clusters in Multi-cell Environments and User and Coordination Pattern Scheduling

[0020] Embodiments of the invention include a method and operation of user scheduling for wireless multi-cell systems in which the pattern specifying the coordination of antennas, or equivalently the “cell” structure with coordinated base-station or distributed antenna systems, or coordination pattern of cells, is made to vary in a controlled fashion over transmission resources. Examples of such resources include slots in time and/or frequency. Such systems are described in more detail in United States Patent Application No. 12/538,729, filed August 10, 2009, titled “A Variable Coordination Pattern Approach for Improving Performance in Multi-Cell or Multi-Antenna Environments”. Such systems are made up of

clusters of base stations coordinating, or cooperating, together for transmissions with user terminals in the respective clusters. The clusters are specified in coordination patterns. In one embodiment, the communication system varies the coordination pattern to enable users to see different interference levels over time (or other transmission resource), where the result is that some levels are more favorable than others. More specifically, coordinating transmissions across multiple cells, and varying such coordination patterns, can help prevent a user from being permanently disadvantaged with a high interference level, as happens to edge users in a system where coordination patterns do not vary.

[0021] Two example coordination patterns that may be used are the patterns illustrated in Figures 2 and 3. Cells of the same shading coordinate their transmissions. That is, the antennas contained within these cells constitute one subset of antennas which coordinate. In Figure 2, users “user4” and “user5” are at the center of a cluster of cells 1, 2 and 3 which coordinate. In contrast, “user6” is at the edge of this cluster. All users within cells 1, 2, and 3 form the user subset served by the antenna subset defined by antennas within these cells. In Figure 2, also, “user7” is at the edge of a cluster of cells 4, 5 and 6 which coordinate. This cluster defines another antenna subset and another user subset. In Figure 3, users “user6” and “user7” are now near the center of a different coordinating cluster of cells made up of cells 3, 4 and 5. Here now other users, e.g. user4 and user5, are at the edge of a cluster made up of cells 1 and 2. Thus, in the system, one can expect user4 and user5 to generally have more favorable interference levels under the pattern of Figure 2 than under the pattern of Figure 3. The opposite is true for users “user6” and “user7”. Using a plurality of such patterns on different transmission resources, e.g. on different time and/or frequency slots (e.g., Figure 2 is used on some frequencies while Figure 3 is used on other frequencies), performance can be greatly increased for all users. This is particularly true for users that were originally “cell edge” users in Figure 1.

[0022] Still, the invention disclosed in United States Patent Application No. 12/538,729, filed August 10, 2009, titled “A Variable Coordination Pattern Approach for Improving Performance in Multi-Cell or Multi-Antenna Environments,” allows for users to be served by multiple patterns. For example, it

is not impossible (or preferable) to limit user6 and user7 to be serviced only when the cluster pattern in Figure 3 is in effect. They may (through random variations in channels or other criteria) also find it advantageous to be serviced for other cluster patterns (e.g., the pattern in Figure 2 or any other pattern). However, for this system to work, given a user has many opportunities over many coordination patterns and transmission resources and instances to be served, there needs to be a system in place to decide which users are scheduled on which cluster patterns, and when scheduled with what rate.

User and Coordination Pattern Scheduling

[0023] To serve all users in a wireless system, and to serve them in an efficient and/or fair fashion, it is often the case that it is either necessary and/or beneficial to decide when to transmit users across the available transmission resources. For example, if resources are in slots in time, in each slot a decision is made on which user or users to serve. Similarly if resources are in bands or tones of frequencies, on each tone or within each band, the system may decide on different subsets of users to serve. This is of course required to share such resources among users in a fashion deemed fair, acceptable, or in accordance to the goal of the system. However, further benefits can be obtained from scheduling. For example, benefits can be gained by taking advantage of what is termed “user diversity”, whereby users are scheduled when it is advantageous to do so (e.g., given variations in channel conditions). This operation is performed by a scheduler executing a scheduling algorithm, and it applies to many wireless systems.

[0024] Embodiments of this invention are focused on a class of wireless systems in which the coordination of antennas, or equivalently the “cell” structure with coordinated base-station or distributed antenna systems, or coordination pattern of cells, of the wireless multi-cell system is made to vary in a controlled fashion over transmission resources. Such resources can be for example defined in terms of time slots and/or frequency bands or tones. Here, with such variation in patterns, the operation and design of an efficient scheduler is not obvious or a trivial extension of systems either with a simple (fixed) single cell structure or a single coordination pattern. In such fixed systems, a given user is only served by one cell,

or only served by one set of coordinated antennas, or only served by one cluster. Opportunities to serve this user are limited to transmission resources available to that single cell, or single coordinated set of antennas, or single cluster. It is therefore sufficient that each cell, coordinated set of antennas (CSoAs), or cluster have its own scheduling algorithm that operates independently of other cells/CSoAs/clusters with respect to this user. Each such algorithm can operate in well known fashions. However, when users can be served by multiple patterns, use of independent algorithms over each cluster over each pattern can become very inefficient. Users are members of multiple cells/CSoAs/clusters, over multiple transmission resources. The system then has multiple pairs of CSoAs and user sets (or multiple pairs of cells and user sets, or multiple pairs of clusters and user sets) where such user sets intersect (have common members). For the sake of simplicity such pairs of sets simple will be referred to herein as “antenna-user pairs”. Balancing when a user is scheduled over each pair requires some coordination among schedulers across cells/CSoAs/clusters and across resources and transmission opportunities.

[0025] The systems described herein include one or more schedulers that execute a scheduling algorithm(s) to perform scheduling for the variable coordination pattern system. Methods and structures for performing a scheduling algorithm that acts in a joint (not independent) fashion over multiple coordination patterns.

[0026] In some embodiments, a single, effective, global scheduling algorithm can be composed of smaller constituent scheduling algorithms that focus on each of the antenna-user associated pairs (or associated cell-user pairs, or associated cluster-user pairs). To operate effectively as a global algorithm, these constituent schedulers pass and share sufficient information between each other so as to operate jointly, or semi-jointly. In some embodiments, such “sufficient information” consists of which users were scheduled and the sum or average rates such users have seen over a given time period. Such sufficient information can allow the joint operation of constituent schedulers to make the same scheduling decisions as a single (global) joint scheduler, implemented by a single processing entity having all necessary information.

[0027] In one embodiment, the way the constituent scheduling algorithms are executed is linked to the way the system varies the coordination pattern of antennas to users. That is, a constituent scheduling algorithm is tied to a given antenna-user associated pair or group, or a pair existing on a given transmission resource.

[0028] In another embodiment, the global scheduling algorithm can determine which coordination patterns to use at a given transmission resource (e.g., at a given time and frequency), and thus directs when and what constituent schedulers operate (e.g., over time and frequency).

[0029] In yet another embodiment, a global central scheduling algorithm, for a given set (or sequence of) transmission resource (e.g., set or sequence of time or frequency instances), examines all users, all antennas, and schedules user terminals according to the constraints of which antennas can signal to which user terminals and when. In this case, a central entity makes the scheduling decisions.

Example of Fair Scheduling

[0030] In one embodiment, a scheduler (whether constituent or global) executes a scheduling algorithm to select which user terminals to transmit to on what transmission resource, and with what rate. For a given resource, the user is a member of an antenna-user pair, which specifies what antennas/cell/cluster serves the user on this resource. The scheduling algorithm ensures that users are given fair access to the transmission resources. Fair is determined with respect to a criterion (e.g., proportional fair rate, maximizing minimum per user rate, etc.). Those skilled in the art are familiar with such algorithms when used for systems in which the antenna-user pair is fixed over all possible resources that may be used to serve a user. It should also be stressed that the criterion under which the scheduler operates can be one of a variety of criteria and do not necessarily have to be “fair criterion”. It can for example be one based on a purposeful differentiated service which tries to give some users more rate than others, for reasons such as cost, revenue, application requirements.

[0031] In order to meet the “fair criterion”, a scheduler keeps track of when users were scheduled, and at what rate, in order to operate in an intended fashion. If

a user is associated with only a single antenna-user pair, it is sufficient that each such antenna-user pair have its own independent scheduling algorithm handling only users in this pair. This for example happens in a fixed cell structure or in cluster systems with a fixed cluster structure. In systems considered in this invention a user is served by different cell/CSoAs /cluster arrangements, and different schedulers. In this case, it is not sufficient that each scheduler keep track only of the rates and users it schedules for a given antenna-user pair.

[0032] As mentioned, embodiments of the invention can use many types of schedulers. In one embodiment, the scheduler tries to schedule users (in a cell or cell cluster) so as to maximize a weighted sum rate criterion

$$\sum_{All\ k} w(k)R(k)$$

where $R(k)$ is the average rate with which user- k is to be (or is) served, and $w(k)$ is a weight, such that $w(k) \geq 0$. Of course, over the operation of the system, the vector of achieved values $R(k)$ is influenced by $w(k)$, the physical layer and the scheduler. Thus these values can additionally be functions that adapt over time, i.e. $R(k)$ may be $R_t(k)$ where “ t ” determines the value at time “ t ”.

[0033] In one embodiment, the choice of the weight is determined off-line. In another embodiment, the weights are selected to make the scheduler give each user a what is termed a proportionally fair share of the rate, proportional to what it can achieve. Such a scheduler is called a Proportionally Fair Scheduler (PFS), which is well-known in the art. For example, see D. Tse and P. Viswanath, Fundamentals of Wireless Communications, Cambridge University Press, Cambridge, 2005. In such an embodiment, to select weights in order to achieve PFS, approximately, one can select weights at the j -th scheduling event as

$$w_j(k) = \frac{1}{\hat{R}_j(k)}$$

where $\hat{R}_j(k)$ is an estimate of the rate a user has received up to the j -th scheduling event (“ j ” being similar to the time “ t ” but in units of scheduling instances). In one embodiment, the estimate is done based on past scheduled events. In one embodiment, the estimate is the average rate received by the user up to that point “ j ”

in time. In another embodiment, the estimate is a weighted average rate (e.g., an exponential weighting going back in time). For a system to operate in the desired fashion, the estimate should reflect the average rate a user sees over all possible resources that are serving that user.

[0034] In a system which serves only one user per antenna-user pair at the j -th scheduling event, a scheduler determines rate it can serve each user, and which user to serve based on the weighted metric, given the weights. For example, assume based on the channel state information or signal to noise ratio, or signal to interference and noise ration, of its users, it is determined that at the j -th scheduling event a rate $Z_j(k)$ can be the rate sent to the user k , if it is scheduled alone. Then the scheduler tries to select the user “ k^* ” so as to maximize the scheduled weighted sum rate at the j -th interval, as determined by the user “ k^* ” maximizing the following equation:

$$k^* = \arg \max_k w_j(k) Z_j(k)$$

In the embodiment, the weights are determined for a user using information about the user’s scheduled rates over all coordination patterns, not just the antenna-user pair of this (or one) scheduler.

[0035] In systems which can serve more than one user at a given transmission opportunity, a decision is made on a subset of users to serve. A subset of user terminals is denoted as Ω . Therefore, Z is now a function of both k and Ω , i.e. $Z_j(k, \Omega)$. In one embodiment, the scheduling algorithm executed by the scheduler tries to select users forming a subset Ω^* so as to maximize at the j -th interval

$$\Omega^* = \arg \max_{\Omega} \sum_{\text{over } k \text{ in subset } \Omega} w_j(k) Z_j(k, \Omega)$$

Note, in a MU-MIMO wireless communication system in particular, the user rates $Z_j(k, \Omega)$ often does depend on the subset of users scheduled for transmission. For example, Linear Zero Forcing (the beams/zero forcing vectors and power allocations to those vectors) or Dirty Paper Coding solutions depend on the subset Ω of users. Such procedures are known to those familiar with the state of the art. In

the embodiment, the weights are determined for a user using information about the users' scheduled rates over all coordination patterns, not just the antenna-user pair of this (or one) scheduler.

[0036] Recall, for a system to operate in the desired fashion, information such as the rate estimate $\hat{R}_j(k)$ should reflect the average rate a user sees over all possible resources that are serving that user. For each associated pair of transmit antennas and users, a naïve and inefficient system can operate scheduling algorithms that are independent of the scheduling algorithms run for other pairs. Here the values “ w_j ” (equivalently values $1/\hat{R}_j(k)$ for schedulers such as PFS) are values local to each scheduling algorithm, and independent (and different) from scheduler to scheduler despite the fact that a user- j may be common to many antenna-user pairs. In this sense, each of the scheduling algorithms operate independently for each pair, making whatever use of the user diversity, transmission options, etc. represented in the pair. This is the standard operation in most prevailing systems since the associated pair of antennas and users is fixed, and a user is a member of only one such pair. When antenna-user pairs change over transmission resources, so that a user is a member of at least two or more pairs over at least two or more resources, independent operation is not efficient.

[0037] The scheduling algorithms for systems of interest in this invention, in which users are served over multiple antenna-user pairs, do not operate independently. In some embodiments, independent schedulers are replaced by “constituent” scheduling algorithms that work together to effectively implement a global scheduling algorithm. To do this, if two different pairs of antenna and users share some common users, the operation of the schedulers for such pairs is made dependent. In one embodiment, this is achieved effectively by changing the calculation of the weights “ w_j ”. Specifically the weight “ w_j ”, for a user j common to two or more antenna-user pairs (common to two or more constituent schedulers) is a value that depends on the operation of all corresponding schedulers that can schedule this user. In one embodiment, there is a global common “ w_j ” value, stored possibly in a central repository, and requested when needed by the constituent schedulers. To make this value meaningful, schedulers would also send updates on

the average rate at which they are serving user j so that the central repository can adjust and update the global common " w_j " value accordingly to reflect the average rate over all schedulers. In one embodiment, there is a global common " w_j " value that is shared via transmissions forwarded between constituent schedulers that can schedule user j .

[0038] In another embodiment, each scheduler has their own copy of " w_j ", where local values may not be the same. However, the values are dependent on each other and are approximately the same. To make values dependent in this way, information about user- j is exchanged between schedulers to ensure the individual values of weight " w_j " are either the same, or sufficiently similar (and thus dependent). For example, in one embodiment, all schedulers send their own local copies of " w_j " to all schedulers that handle user j . Each scheduler replaces their local copy of " w_j " by a (sufficiently) weighted mean of all the individual " w_j " values it receives from other schedulers. The weighting of such a mean may reflect how much resources each scheduler has at its disposal, where more weighting is given to " w_j " values from schedulers with more resources. Then each scheduler would operate for a number of transmission opportunities using this value, and operating independently of other schedulers, updating its local measure of " w_j " based on its own independent scheduling events, after which local copies of " w_j " are exchanged again, and the process repeats. Thus individual weight values "track" each other and are sufficiently close.

[0039] An interesting feature of such a system is that the schedulers do not have to exchange information about each and every scheduling event, or interact with schedulers about each and every scheduling event. They (for many schedulers including PFS) sufficiently only need to exchange information such as the sum of rates with which they serve a user over intervals of time.

Deployment and Inter-connectivity of Constituent Schedulers

[0040] In one embodiment, the scheduling system is composed of a number of constituent schedulers. Figure 6 is a diagram of a communication system that includes constituent schedulers that serve common users and share information (such as that mentioned in [0038-39] about those users).

[0041] Referring to Figure 6, the system is shown with two coordination patterns, which are depicted in Figures 4 and 5. Figure 4 illustrates a number of such schedulers S1, S2, S3, S4, corresponding to the clusters of Figure 2. Scheduler S1 is the scheduler for base stations (BS) 1, 2 and 3 of cells 1, 2 and 3 respectively, scheduler S2 is the scheduler for base stations 4, 5, and 6 of cells 4, 5 and 6 respectively, scheduler S3 is the scheduler for base stations 7 and 8 of cells 7 and 8 respectively, and scheduler S4 is the scheduler for base station 10 of cell 10. Figure 5 illustrates a number of such schedulers S5, S6, S7, S8, corresponding to the clusters of Figure 3. Scheduler S5 is the scheduler for base stations 3, 4, and 5 of cells 3, 4 and 5 respectively, scheduler S6 is the scheduler for base stations 1 and 2 of cells 1 and 2 respectively, scheduler S7 is the scheduler for base stations 6 and 10 of cells 6 and 10 respectively, and scheduler S8 is the scheduler for base station 8 of cell 8. Connections between schedulers and base-stations signify which base-stations are used in the respective cluster of the scheduler. Note that connections can be via wired or wireless communication.

[0042] As shown together by Figures 4 and 5, user terminals can be members of different clusters over time or frequency. For example, user5 and user6 are scheduled by scheduler S1 in the channel resource used (e.g., frequency and/or time interval) by Figure 4, and by scheduler S5 in channel resource used (e.g., frequency and/or time interval) by Figure 5. This reflects operations of the varying multi-cell systems of interest.

[0043] Referring back to Figure 6, each of the schedulers receive from the physical layer channel state information that provides information on the state of the channels that exist between user terminals in a cluster and antennas in a cluster. The schedulers also know which user terminals are associated with each of the antennas. In one embodiment, the schedulers also have access to information about one or more of the users' estimated service rate $\hat{R}_j(k)$, users' rate requirements, users' previous scheduled events, and users' assigned rates.

[0044] In this embodiment, the schedulers behave in a joint fashion, and different schedulers handle the same users over different transmission resources, e.g. over different time and/or frequency intervals. They do this by sufficiently sharing information that is sufficient for scheduling. The sharing of such

information is illustrated by the lines of communication between schedulers in Figure 6. More specifically, schedulers S1 and S6 share information about common users in cells 1 and 2, schedulers S8 and S9 share information about common users in cell 3, schedulers S2 and S5 share information about common users in cells 4 and 5, schedulers S2 and S7 share information about common users in cell 6, scheduler S2 and S8 share information about common users in cell 8, and schedulers S4 and S7 share information about common users in cell 10. For example, for a Proportionally Fair Scheduler (PFS), they share sufficient information to make " w_j " values of common users either the same or approximately the same.

[0045] In one embodiment, the scheduler determines the rate information by sharing only sufficient information for this purpose by directly transferring such information between constituent schedulers (without use of a central repository). In one embodiment, the only information that is shared is the sum or average rates of such users are given by each of the schedulers. For example, in one embodiment, it is sufficient for S1 to tell S5, and S5 to tell S1, when and with what rate it has scheduled user terminals such as user5 and user6 common to both schedulers. In other embodiments, this may be more than sufficient information. For example, if $w_j(k)$ is set as the average rate as discussed above, it is also sufficient for S1 to tell S5, and for S5 to tell S1, only the average rate that it has served such users over a given number of transmission opportunities. Therefore, each scheduler can, with the average for each coordination pattern, weigh each average over the number of scheduling opportunities a scheduler has, and compute the average rate each user terminal has been served over all coordination patterns.

[0046] It may also be sufficient for S1 to tell S5, and S5 to tell S1, the weight it used. In one embodiment, weights are updated within each scheduler (e.g., S1) based on its scheduling of user terminals, and the change in such weights as seen from interval to interval by other schedulers (e.g., S5) provides information as to how a user terminal was served.

[0047] In one embodiment, the weights are determined for a given user terminal using past information on when and at what rates user terminals were scheduled. In one embodiment such weights are determined for a user terminal using information about the user terminal's scheduled rates over all coordination

patterns (all schedulers which consider the user). At the same time, when a constituent scheduler is making a decision (e.g., scheduler S1), it is sufficient for scheduler to know or determine only the weights " w_j " of its user terminals. For PFS, this is just the inverse of the average rate.

[0048] In one embodiment, the schedulers also ensure user terminals are scheduled with a maximum latency. Such a maximum latency value may be application dependent. For example, in voice communication, it is necessary to send a packet at least once every block of x msec, where " x " depends on the system but is typically on the order of 10 to 30 msec. If schedulers also have some requirement to ensure users are scheduled (given some rate) with some maximum latency, it is sufficient for S1 to tell S5, and S5 to tell S1, only the average rate (or weights) that it has served such user terminals, and the last time it served such user terminals.

[0049] In one embodiment, weights are determined by some other criteria (e.g., how much a user pays, the application it is running, etc.). If the weights are determined by some other criterion/criteria, the scheduler may not in fact update weights, but would use (or set) weights according to that criterion/criteria.

[0050] In one embodiment, the channel state information is not (necessarily) exchanged between the constituent schedulers. The information shared is sufficiently limited to only the scheduling information (i.e., information sufficient for scheduling). In one embodiment, each constituent scheduler uses the channel state information to compute the potential scheduled rates $Z_j(k)$, along with the computed weights, that apply only to users in its cluster. Another constituent scheduler does not need to know such channel state information. It just needs to know the rates (or sum rate) for each user scheduled.

[0051] In another embodiment, which can use similar scheduling techniques and information sharing requirements as described above, the information used by the schedulers to operate in joint or semi-joint fashion is stored in one or more central repository locations. In one embodiment, for PFS schedulers, the information in the central repository should be sufficient to provide a measure of the average rate a user has received up to a given time over all possible transmission resources it has. Such information thus does not need to include each and every

scheduling instance, and the rates for each instance; it only needs to be sufficient to determine the average rate a user has seen. Sufficient information for such a purpose could be periodic (not real time) updates on the total rate each constituent scheduler has given the user. As such, a total rate over all schedulers can be determined and converted to an average measure over various intervals in time.

[0052] Figure 7 illustrates the case with a single central location at central entity 701 serving all of the scheduling entities S1-S8. In one embodiment, this central entity 701 resides for example on a backbone network (e.g., an IP based network). The latencies over such a backbone may be large compared to the interval of scheduling events. In one embodiment, if this is true, the system considers that information up to a certain time is sufficient to do the scheduling in the j -th interval. For example, it is possible that the central entity 701 may only have, at the j -th interval, information up to the $(j-K)$ -th interval, where K is some acceptable number. In this way, coordination of the scheduling over the entire system does not have to be done with the stringent latency requirements as may be required by a full coordination at the physical layer signaling.

[0053] The schedulers S1-S8 request updates from central entity 701 either before every slot they operate, or with some regular interval. The schedulers S1-S8 entities also send updates to central entity 701 either after every slot they operate, or with some regular interval. When updating over intervals, often information need only reflect the sum or average of such information over the interval, e.g. the sum or average rate with which a scheduler served each of its users over the interval.

[0054] Scheduling entities may also send information via other entities. As shown in Figure 7, scheduler S6 sends information through scheduler S1, scheduler S5 sends information through scheduler S3, and scheduler S7 sends information through scheduler S8. This allows

1. Some entities to get necessary information they need before it is sent to the central entity. In fact, they would therefore not need to request such information from a central repository.
2. Possibly to simplify the backbone interconnection of entities.
3. Reduce the information requested from the central entity, particularly where all information comes between entities. For example, they may be able to

obtain all the information they need for their users through such inter-scheduler exchanges and forwarding of information.

[0055] In another embodiment, the information sent from the central entity to schedulers is only the weights $w_j(k)$. The central entity maybe able to calculate such weights. This calculation may be made by seeing a sufficient amount of past scheduling information from all schedulers, and then sending only the weights needed to each of the constituent schedulers.

[0056] In yet another embodiment, only a subset of controllers needs to know and track the rates of any given user terminal for scheduling. Consider an arbitrary user terminal k . Let $A(k)$ denote the set of controller/schedulers that can schedule user k . In one embodiment, another set of entities B are associated with storing scheduling information about distinct user terminals. In this embodiment, one or more entities in the union of the sets $A(k)$ and B is used as scheduling information repository for user terminal k . As such, the repositories for the “ k ”th user terminal are responsible for providing the scheduling algorithm weights for the “ k ”th user terminal to the controllers scheduling that user terminal at an given time, as well as for knowing the scheduled rates of the “ k ”th user terminal and calculating its updated weights.

[0057] In one embodiment, the repositories are chosen from user terminal k from the set $A(k)$ only. The repositories can be selected according to, for example, one of the following methods: (i) the controller that is closest to user terminal k is chosen as a single repository for that user; (ii) each of the controllers in the set $A(k)$ is chosen as repository for a given period (e.g., a round robin fashion over some period of transmission opportunities); (iii) all the controllers are chosen as repositories; and (iv) a subset of controllers is a priori designated as repositories and all the designated repositories in the set $A(k)$ is chosen as repositories for user k .

[0058] Figure 8 is a flow diagram of one embodiment of a scheduling process. The process is performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. The scheduling process is performed by a scheduler.

[0059] Referring to Figure 8, the process begins by processing logic receiving information on which coordination pattern to use (processing block 801). This pattern could be predetermined, e.g. as determined by a sequence of stored pattern decisions. After receiving information on which coordination pattern to use, for each coordination cluster (antenna-user pair) in the pattern, processing logic in the constituent scheduler serving this antenna-user pair requests and obtains information (if it does not already have it) from a central repository or other schedulers (processing block 802) and creates necessary information to make local scheduling decisions (processing block 803). In one embodiment, the necessary information includes information to determine the weights in a PFS algorithm.

[0060] Next, for each coordination cluster in the pattern, processing logic obtains an estimate of the rate it can support to a user terminal or terminals (processing block 804). In one embodiment, the estimates include estimates of user terminal SINRs from channel state information. In another embodiment, the estimates include estimates of rates supported for various subsets of users. In yet another embodiment, the estimates include both the user terminal SINRs and the rates supported for the various subsets of users.

[0061] Then, for each coordination cluster in the pattern, processing logic selects a user terminal, or subset of user terminals to be scheduled (processing block 805), updates and stores information it and other schedulers require for future scheduling (processing block 806), and sends the information (required to update weights or other variables for common users) to a central repository, if necessary (processing block 807). Such updates can happen after every scheduling event, or after a number of scheduling events as previously described. Thereafter, the process transitions to processing block 801 and the process repeats.

[0062] In another embodiment, a central repository is not used. Here processing block 802 does not get information from a central repository, but possibly from a number of repositories and/or a number of other constituent schedulers as, for example, previously described when discussing Figure 6.

[0063] In another embodiment, processing block 807 does send update information to a central repository, but possibly to a number of repositories and/or a

number of other constituent schedulers as previously described when discussing Figure 6.

A Coordination Controller Embodiment

[0064] In one embodiment, the communication system is as shown in Figure 6, includes a single central entity and an additional entity referred to herein as a coordination controller. This is illustrated in Figure 9. A flow diagram of another embodiment of a scheduling process used in such a system is shown in Figure 10.

[0065] Referring to Figure 9, the system includes coordination controller 901. Coordination controller 901 has access to the centrally collected scheduling information 902, which is stored in memory accessible via wired or wireless connections. If coordination controller 901 operates with schedulers that are PFS, then coordination controller 901 can also poll constituent entities directly for updates on such information (e.g. weights, rates scheduled in the past, etc.).

[0066] In one embodiment, controller 901 determines, given scheduling history, if it is advantageous to change the operation of the coordination patterns. For example, if it is noted that some user terminals are getting a higher rate than needed, and this happens when certain coordination patterns operate, while other users are not getting what they require, but are better served by other coordination patterns, the controller may decide to change the relative number of transmission opportunities or assigned transmission resources (e.g., time and/or frequency slots) assigned to various patterns. For example, if user6 and user7 have much higher rate requirements than user4 and user5, and user terminals such as user1 and user2 get the rate they need regardless of which pattern is used, then the controller may decide to use the pattern in Figure 3 more than the pattern in Figure 2. This could have advantages in cases where the number of patterns is limited, or where rate requirements are set externally, outside of the PFS, that are not possible with an equal or fixed sharing of patterns.

[0067] In one embodiment, controller 901 also decides to remove or add some coordination patterns into the system as required.

[0068] Figure 10 is a flow diagram of one embodiment of a scheduling process. The process is performed by processing logic that may comprise hardware

(circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. The scheduling process is performed by a scheduler.

[0069] Referring to Figure 10, the process begins by processing logic decides on a coordination pattern (processing block 1001). After receiving information on which coordination pattern to use, for each coordination cluster in the pattern, processing logic in the scheduler requests and obtains information from a central repository or other schedulers (processing block 1002) and creates necessary information to make local scheduling decisions (processing block 1003). In one embodiment, the necessary information includes any such information that is able to specify the weights in a PFS algorithm.

[0070] Next, for each coordination cluster in the pattern, processing logic obtains an estimate of the rate it can support to a user terminal or terminals (processing block 1004). In one embodiment, the estimates include estimates of user terminal SINRs from channel state information. In another embodiment, the estimates include estimates of rates supported for various subsets of users. In yet another embodiment, the estimates include both the user terminal SINRs and the rates supported for the various subsets of users.

[0071] Then, for each coordination cluster in the pattern, processing logic selects a user terminal, or subset of user terminals to be scheduled (processing block 1005), updates and stores information it and other schedulers require for future scheduling (processing block 1006), and sends the information to a central repository, if necessary (processing block 1007). Thereafter, the process transitions to processing block 1001 and the process repeats.

Constituent Schedulers for Multiple Clusters

[0072] This embodiment shares principles of the previous embodiments. The difference is that a single entity is used to schedule over multiple clusters. An example of this is shown in Figure 11.

[0073] Referring to Figure 11, a single entity CS1 is used to schedule over two clusters taken from Figure 2 at a given time (or frequency). In one embodiment, this scheduler also coordinates over another pattern, at a different time

or in a different frequency band, for example the cluster represented by Cell 3 + Cell 4 + Cell 5 in Figure 3. Therefore, it may be able to make concurrent joint decisions on user terminals common to both time or frequency intervals.

[0074] Also, in one embodiment, a scheduler such as CS1 makes concurrent joint decisions on scheduling users and scheduling which coordination patterns to user. In one embodiment, if Figure 2 operates in one frequency band, and Figure 3 in another, at the same time, it can consider the scheduling of user6 and user7 jointly.

[0075] In one embodiment, the scheduler CS1 decides to schedule in each band in a way such that a user terminal is not scheduled over both (or more than a given number of) bands simultaneously. In another embodiment, the scheduler CS1 decides to schedule over bands jointly which enable further cooperative signaling over coordination patterns in different frequencies. For example, in an OFDM system the signal may be split over multiple bands. However, in forming the signal one needs to know ahead of time the supported rates and how many bands are being used.

[0076] The interconnectivity represented in Figure 11 also has implications to the amount of information that needs to be passed between these wider entities. In one embodiment, there may be only one entity. For example, the coordination controller in Figure 9 only does the scheduling, in which case “S1”, ..., “S8” may be less complex entities.

Further Design Considerations

Physical Layer and Pattern Considerations In One Embodiment

[0077] In the above-described embodiments, the scheduling algorithm determines, based on information (e.g., channel state information), the rate that a user can be scheduled at a given point in time. The rates over all users are compared, and a subset of users are scheduled in line with the over all criterion (e.g., a weighted criterion). The embodiments do not necessarily depend on the physical layer. In a multi-cell system, the physical layer transmissions within each cluster of cells operate independently of other clusters. For purposes herein the transmitting signals themselves can use any technique it chooses, e.g. single-input

single-output (SISO) transmission; multiple input multiple output (MIMO) transmission; and, multi-user MIMO (MU-MIMO) transmission whereby multiple antennas coordinate a joint concurrent transmission to multiple users. Underlying transmissions can be based on Orthogonal Frequency Division Multiplexing (OFDM), Code Division Multiple Access (CDMA), etc.

[0078] MU-MIMO is of particular interest in these multi-cell designs, and has nice properties for scheduling algorithms since constituent schedulers can in fact assume that interference within each cluster is controlled. By jointly coordinating transmissions over every cell within a cluster, using MU-MIMO as the underlying signaling, these can be significant benefit in alleviating problems due to interference. MU-MIMO transmission in the form of Linear Zero Forcing (LZF), Dirty Paper Coding (DPC), or Minimum Mean Square Error (MMSE) driven signaling, all have the advantage of controlling the effect of interference within each cluster. When MU-MIMO is used with PFS, one can look at various techniques to jointly select the user terminals and the MU-MIMO physical layer signaling solution.

[0079] With Linear Zero Forcing (LZF), one can consider a greedy scheduling algorithm which starts in step 1) by selecting only one user terminal in the cluster based on the assumption it is scheduled alone (in which case the LZF solution is trivial), then searches for the best combination of 2 users, limiting the search such that the user in step 1) is one of the users (the LZF solution changes with 2 users), and so on until a maximum number of user terminals is selected (the LZF solution changes with each additional user terminal).

[0080] In one embodiment, the selection of user terminals at each step would look at individual and sum weighted rates $w_j(k)Z_j(k)$ and sum over potential users $w_j(k)Z_j(k)$ in making decisions.

Physical Layer and Pattern Considerations in Another Embodiment

[0081] In one embodiment, the scheduling algorithm includes in the rate terms $Z_j(k)$ provision for interference coming from outside of the cluster.

[0082] In one embodiment, in Figure 2, one could provide estimates of average interference levels seen by different users. This information can be

estimated by the users themselves, by sensing transmission levels they see from stations which do not serve them. Such information can be implicit on a user's geographic locations. In one embodiment, such an estimate is derived based on geographic considerations and assumed powers radiating from antennas from other cluster(s). For example, user6 would probably be given a different (higher) estimated (whether sensed or determined by location) interference than users1, 2, 3, and 4 for this coordination pattern

[0083] Furthermore, since users are members of different clusters at different times, in one embodiment, they in fact also have partial information useful in helping to give estimated interferences. For example, in Figure 2, user6 is not being served by transmissions from base stations BS4 or BS5. However, when the system is operating as in Figure 3, it does have information on the channel between itself and BS4 and BS5. In one embodiment this information is used in determining the interference under the operation of Figure 2. This would for example assume that, under Figure 2, BS4 and BS5 are radiating about the same power/signal as it did under Figure 3.

Other Physical Layer and Pattern Considerations

[0084] The embodiments of the joint scheduling and coordination pattern variation can use the concepts described in other pattern varying schemes such as those in which, power levels applied to different frequency bands over the clusters over the coordination patterns can be chosen in a way which helps the scheduling. See United States Patent Application No. 12/538,739, filed August 10, 2009, titled "A Method for Varying Transmit Power Patterns in a Multi-Cell Environment".

Objectives with Respect to Coordination Patterns and the Users They Serve

[0085] User terminals such as user1 and user2 in Figure 2 may overwhelm other user terminals, since the rates they may be able to be served with are good irrespective of the coordination pattern used. As a result, it is possible that the scheduler, unless it reacts to this imbalance, may not serve users such as user4, user5, user6, or user7 sufficiently. This is true even when using algorithms such as PFS.

[0086] As an extension to embodiments, and particularly as an extension to some of the concepts in described above, one can consider schedulers which further limit the subset of user terminals being scheduled at each constituent scheduler, or within each coordinated cluster, in a purposeful fashion. In one embodiment this can, for example, be done as a function of frequency band depending on either the user's geographic location or a user's estimated rate. For example, in Figure 2, a scheduler reserves some bands for user terminals such as user4 and user5, prohibiting user terminals such as user1 and user2 from being scheduled in such bands.

[0087] In one embodiment, one purposefully changes the weighting in the scheduling algorithm to accomplish this. Restricting a user terminal is equivalent to setting its weight $w(k)$ to zero. In one embodiment, discouraging the scheduling of a user terminal is accomplished by reducing $w(k)$ further (e.g., scaling $w(k)$ and using a new value $p(k)w(k)$, where $0 \leq p(k) \leq 1.0$). This scaling can be used when scheduling in a given band or a given coordination patterns, and would over-ride some of the default operation of a PFS.

Advantages of Embodiments of the Invention

[0088] Embodiments of the invention include a joint user+rate+pattern scheduling algorithm within a multi-cell environment for which the coordination pattern between cells vary. Furthermore, the scheduling itself can determine the scheduling of coordination patterns.

[0089] Embodiments of the invention have one or more of the following advantages. The joint operation allows the scheduling algorithm, which is made of many smaller "constituent" scheduling algorithms operating in a joint fashion, to make better decisions about which user terminals to schedule and when. By allowing the scheduling algorithm effectively link all possibilities (all possible scheduling opportunities) over all coordination patterns (all possible antenna-user pairs over all transmission resources), user terminals can be scheduled when it is best with respect to all such possibilities. This has advantages in terms of the total rate served to user terminals, with respect to fairness between users, with respect to individual rates user terminals see, and with respect to latency of serving user

terminals. The joint operation also allows the scheduling algorithm to influence the operation of the variation of coordination patterns. The scheduling algorithm may in fact also direct this operation, making a fully joint decision not only over all constituent schedulers, but over the scheduling of patterns themselves.

An Example of a Scheduler

[0090] Figure 12 is a block diagram of one embodiment of a scheduler that performs one or more of the functions described above, including performing a scheduling algorithm for a cluster or for multiple clusters. In one embodiment, schedulers are located within base stations. In another embodiment, the scheduler is separate from base stations.

[0091] Referring to Figure 12, the scheduler 1200 includes a

[0092] a control interface 1202 to send and receive control information. It can use this interface to send information to control base stations, such as, for example, information specifying which user terminal to schedule in which cluster and on which transmission resource, as well as transmission rates for the user terminals.. It can also use this information to receive information required to make its scheduling decisions, such as user weights " w_j ", user average or sum rates from other schedulers and or a central repository, user channel state information, user estimated SINRs, etc. The control interface 1202 may interface with one base station or multiple base stations and may comprises wired or wireless communication functionality.

[0093] Processor(s) 1201 executes one or more scheduling algorithms, such as those described above, to control, or cause to be controlled, many of a base station's operations (or operations of other base stations). This would include selecting user sets such as Ω given the scheduling criterion and information received from to control interface 1202. User rates are also determined by this unit. It can make decisions on what coordination pattern to use. Memory 1203 stores instructions and other data, along with pre-defined coordination patterns and scheduling algorithm(s), such as those described above. A timing/clock unit 1204 provides timing and/or clocking for the base station, in a manner well-known in the art.

[0094] Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no way intended to be considered limiting. Therefore, references to details of various embodiments are not intended to limit the scope of the claims which in themselves recite only those features regarded as essential to the invention.

CLAIMS

We claim:

1. A wireless communication system comprising:
a plurality of user terminals;
a plurality of antennas to communicate with the plurality of user terminals, wherein the plurality of antennas are grouped into coordination patterns that specify which antennas operate together as a cluster to coordinate transmissions with a subset of the plurality of user terminals, and further wherein the defined pairs of the antenna sets and associated user sets vary over transmission resources based on coordination pattern changes such that different sets of antennas in the plurality of antennas are grouped together at different times to coordinate transmissions to different subsets of users; and
at least one controller to execute at least one scheduling algorithm to schedule transmission with user terminals by specifying which user terminal to schedule for service by which set of antennas and on which transmission resource and with what transmission rate.
2. A method comprising:
varying clusters of antennas into new coordination patterns that specify which antennas operate together as a cluster to coordinate transmissions with a subset of the plurality of user terminals, wherein clusters vary over transmission resources based on coordination pattern changes such that different sets of antennas in the plurality of antennas are grouped together at different times to coordinate transmissions to different subsets of users; and
executing at least one scheduling algorithm to schedule transmissions with user terminals by specifying which user terminal to schedule in which cluster and on which transmission resource and transmission rates for the user terminals.
3. A method for scheduling user terminals in a wireless communication system having a plurality of user terminals and a plurality of

antennas that are grouped based on coordination patterns that specify which antennas operate together as a cluster to coordinate transmissions with a subset of the plurality of user terminals, the method comprising:

for each coordination cluster specified in a coordination pattern, performing operations including

receiving information required to implement a scheduling criterion;

creating information to make a scheduling decision;

obtaining an estimate of a rate the system using the set of antennas can support to a user terminal;

selecting the user terminal for scheduling based on the rate estimate;

determining the rate a user is to be given if it is scheduled.




FIG. 1

Cells of the same color (shading) coordinate transmissions




FIG. 2

3/12

Cells of the same color (shading) coordinate transmissions




FIG. 3

4/12

Cells of the same color (shading) coordinate transmissions




FIG. 4

5/12

Cells of the same color (shading) coordinate transmissions




FIG. 5

6/12

SCHEDULERS CONNECTED BY DASHED LINES SHARE INFORMATION IN DOING SCHEDULING
Cells of the same color (shading) coordinate transmissions




FIG. 6

7/12

Cells of the same color (shading) coordinate transmissions




FIG. 7

8/12




FIG. 8

9/12




FIG. 9

10/12




FIG. 10

11/12




FIG. 11

12/12




FIG. 12

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2009/053466

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04W16/28

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04W H04B

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	PAPADOGIANNIS A ET AL: "A Dynamic Clustering Approach in Wireless Networks with Multi-Cell Cooperative Processing" COMMUNICATIONS, 2008. ICC '08. IEEE INTERNATIONAL CONFERENCE ON, IEEE, PISCATAWAY, NJ, USA, 19 May 2008 (2008-05-19), pages 4033-4037, XP031266078 ISBN: 978-1-4244-2075-9 page 4033, right-hand column, line 30 - line 35 page 4034, left-hand column, line 10 - line 22 page 4034, left-hand column, line 25 - line 36 page 4035, right-hand column, line 14 - line 26 page 4036, left-hand column, line 16 - line 47 -/--	1-3

☒ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

9 November 2009

Date of mailing of the international search report

18/11/2009

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040.
Fax: (+31-70) 340-3016

Authorized officer

López García, J

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2009/053466

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	page 4036, right-hand column, line 3 - line 11 ----- KARAKAYALI M K ET AL: "Network Coordination for Spectrally Efficient Communications in Cellular Systems" IEEE WIRELESS COMMUNICATIONS, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 13, no. 4, 1 August 2006 (2006-08-01), pages 56-61, XP001547066 ISSN: 1536-1284 page 56, left-hand column, line 16 - right-hand column, line 31 page 57, left-hand column, line 49 - page 58, right-hand column, line 32	1-3
A	----- ETRI: "Coordinated multi-cell transmission for LTE-Advanced downlink" 3GPP DRAFT; R1-082896 MULTI-CELL MIMO, 3RD GENERATION PARTNERSHIP PROJECT (3GPP), MOBILE COMPETENCE CENTRE ; 650, ROUTE DES LUCIOLES ; F-06921 SOPHIA-ANTIPOLIS CEDEX ; FRANCE, no. Jeju; 20080812, 12 August 2008 (2008-08-12), XP050316374 [retrieved on 2008-08-12] page 1, paragraph 2.1 - page 2	1-3
A	----- SAMSUNG: "Inter-Cell Interference Mitigation Through Limited Coordination" 3GPP DRAFT; R1-082886 INTER-CELL INTERFERENCE MITIGATION THROUGH LIMITED COORDINATION, 3RD GENERATION PARTNERSHIP PROJECT (3GPP), MOBILE COMPETENCE CENTRE ; 650, ROUTE DES LUCIOLES ; F-06921 SOPHIA-ANTIPOLIS CEDEX ; FRANCE, no. Jeju; 20080812, 12 August 2008 (2008-08-12), XP050316366 [retrieved on 2008-08-12] page 1, paragraph 1 - page 2 page 4, paragraph 4.1 - page 6 -----	1-3