June 20, 2018

ACTIVE DIRECTORY AND CALENDARING BASED ALLOCATION AND RESERVATION OF WIRELESS RESOURCES

Sukrit Dasgupta

Jp Vasseur

Santosh Pandey

ACTIVE DIRECTORY AND CALENDARING BASED ALLOCATION AND RESERVATION OF WIRELESS RESOURCES

AUTHORS:
Sukrit Dasgupta
Jp Vasseur
Santosh Pandey

ABSTRACT

Present techniques dedicate wireless channel resources based on applications that will be used in collaborative meetings. The degree of reservation is determined using machine learning algorithms that have knowledge from a prior history of meetings, participants in the meeting, etc. This ensures that collaborative applications will perform favorably, leading to more effective collaborations. Several novel components are involved in these techniques, which control the network infrastructure, determine the users that need reservation(s), identify the applications that will be used, and also utilize machine learning to model the usage and requirements.

DETAILED DESCRIPTION

In any enterprise, there are events where a large number of client devices aggregate in larger than usual numbers in certain geographical areas. Examples of such locations include community rooms, meeting rooms, etc. and the meetings are often planned in advance. During these meetings, one or more specific applications are used for collaboration, information exchange, or other form of communication. Due to limited and fixed wireless infrastructure (e.g., hardware, channel resources, etc.) in these locations, user experience is often degraded as resource requirements are higher than normal. Techniques are provided for resource reservation *a priori* and to move surrounding users to alternative different access points (APs), or alternatively, to create a dedicated service set identifier (SSID) on nearby APs that support only this particular event, and thus, prevent other users from joining the network through these APs.

The present techniques include several components that together provide the functionalities and responsibilities described herein. The calendar registry component

Published by Technical Disclosure Commons, 2018

registers with the calendaring service of the enterprise to receive continuous updates of new meetings being scheduled. In addition, the calendar registry component receives details of the users, locations, type of media, and time duration of the reservations being made. In some aspects, the calendar registry directly interacts with the calendaring service as well as the network coordination module.

Based on this information, the system determines the level of required resources (learned from past events) to determine the expected traffic at the time of the upcoming event being scheduled by the calendar registry. The first step involves determining if the planned event is a recurring event. This may be accomplished by enabling machine learning (ML) seasonality detection on the event itself. If the same event is planned at the same time every 'X' time/days/months/years/etc., a seasonality detection module can be used, e.g., a forecasting tool which performs seasonal decomposition, or a logistic regression using time/day/months/years/etc. as input feature(s). First, the system determines whether the event is recurring or whether the event is "exceptional". In the second case, the extra resource (ER) required can be added to the level of usual resources (UR) typically observed at that location on the network (e.g., the average level, the 90th percentile level, the maximum level, etc.). Thus, the expected level of resources is ER+UR. In some cases, the level of resources can be expressed in terms of the number of clients expected on the network. In other embodiments, the level of resources can account for the type of application expected and the expected traffic load per class of services (CoS), etc. For example, if the present techniques are used exclusively to reserve capacity in the network for a given traffic CoS, resources may be expressed as the number of expected clients using the application for that given CoS.

The network coordination module is responsible for reserving channel capacities in the wireless infrastructure. After receiving the meeting details from the calendar registry component and predicting the expected resources needed, the network coordination module queries the active directory infrastructure to determine the exact location of the meeting and the device types of the users. It then correlates this information with the geolocations of the APs in the network to determine the possible APs that will be involved in servicing the upcoming meeting. This module then sends a reservation request message to the wireless LAN controller (WLC) managing these APs with the requirements, e.g., expressed in terms of required wireless resources, such as ER+UR, or K * (ER+UR) where K is the under/over-booking factor. In yet another embodiment, the factor may be specified on a per-CoS basis. The WLC then relays this message to the corresponding APs.

The channel reservation component provides signaling by the APs to reserve the channel and configure the quality of service (QoS) for the media types predicted to be involved in the meeting. The channel reservation may be translated to a new SSID created on subset of APs while removing other SSIDs. In other cases, it may involve radio resource management (RRM) to assign 40 MHz, 80 MHz, or 160 MHz bandwidth to a few select APs while moving nearby APs to other channels. This allows for a larger capacity and data rate to be supported on these APs, thereby allowing higher traffic loads to be handled. In this case, closed loop control may effectively modify wireless network topology on-the-fly. Furthermore, enhanced distributed channel access (EDCA) parameters on these APs may be modified to be more aggressive for certain traffic, based on applications determined to be critical during the meeting (per CoS). Should the request not be accommodating because of lack of capacity (e.g., two events scheduled in the same area, with one or both requiring too much capacity), a message may be sent back to the calendar registry to indicate that the network will not be able to guarantee a minimum service level agreement (SLA) for the event participant, in view of the extra resources being requested and the usual traffic load observed for this location. Such feedback may prompt the event organizer to relocate the meeting. In yet another embodiment, the system may propose an alternate location where such request may be satisfied, or the system may use a trial and error approach to identify a place (nearby in the same building, in a neighboring building) where such request may be accommodated.

Another component is responsible for monitoring the usage during the duration of the event and learning the channel characteristics so that future requests can be handled more optimally (also referred to as implicit feedback). To achieve this, the system uses telemetry about users, channel quality, interference, application types, TCP control, application mean opinion scores (MoS), etc. In another embodiment, user-based feedback may be gathered from the users once the meeting has concluded. This module gathers information about their experience, quality of voice/video, etc.

The final component detects correlations with the number of issues raised by Anomaly Detection (AD) systems. For example, consider the case where an AD system is deployed in the network, such as DNA Analytics. In this case, it is possible to determine whether anomalies are detected during the event even if the "a priori resource reservation" approach described herein has been enabled. A feed-back signal may be sent to the network coordination module to adjust the resource reservation strategy. For example, if the quality of experience is not satisfying (presence of anomalies), the parameters related to reservation may be increased (e.g., K). Such parameter(s) may also be reduced to minimize the amount of resources that are reserved. In yet another embodiment, if a root cause of the anomaly is discovered, this may be provided to the network coordination component for further action. If an event takes place, resources are reserved by the resource reservation module (RRM) in the wireless space, and anomalies related to poor quality of experience are detected, wherein the root cause lies in the WAN domain (as opposed to the wireless network). Such information may be advantageously used to trigger an action in the WAN domain (e.g., change of QoS on the WAN, set up of a dedicated MPLS traffic engineering tunnel, etc.).

The feedback component uses a reinforcement learning component having implicit and explicit feedback provided to the system to adjust the parameters (UR, ER and K). Once the meeting duration expires, this component is also responsible for freeing channel resources for general usage and reactivating the closed loop control, which is continually used to improve the model that predicts and allocates channel resources for future events.

In summary, techniques are provided to dedicate wireless channel resources predicted to be needed for upcoming collaborative meetings. The degree of reservation is determined using ML algorithms that have knowledge from a prior history of meetings, participants in the meeting, etc. Various components control the network infrastructure and reservation process as well as determine the number of users that need reservation(s) and corresponding user applications. The ML component may model usage and requirements to improve performance in order to make collaborations more effective.