

Kenjiro Cho
Philippe Jacquet (Eds.)

LNCS 4311

Technologies for Advanced Heterogeneous Networks II

Second Asian Internet Engineering Conference, AINTEC 2006
Pathumthani, Thailand, November 2006
Proceedings




Springer

Kenjiro Cho Philippe Jacquet (Eds.)

Technologies for Advanced Heterogeneous Networks II

Second Asian Internet Engineering Conference, AINTEC 2006
Pathumthani, Thailand, November 28-30, 2006
Proceedings



Volume Editors

Kenjiro Cho

Internet Initiative Japan, Inc.

1-105 Kanda Jinbo-cho, Chiyoda-ku, Tokyo 1010051, Japan

E-mail: kjc@ijjlab.net

Philippe Jacquet

Institut National de Recherche en Informatique et Automatique

Campus of Rocquencourt, Domaine de Voluceau, BP105, 78153 Le Chesnay, France

E-mail: Philippe.Jacquet@inria.fr

Library of Congress Control Number: 2006936488

CR Subject Classification (1998): C.2.4, C.2, C.3, F.1, F.2.2, K.6

LNCS Sublibrary: SL 5 – Computer Communication Networks and
Telecommunications

ISSN 0302-9743

ISBN-10 3-540-49364-6 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-49364-8 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006

Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper SPIN: 11930181 06/3142 5 4 3 2 1 0

Preface

The Asian Internet Engineering Conference (AINTEC) brings together researchers and engineers interested in practical and theoretical problems in the Internet technologies. The conference aims at addressing issues pertinent to the Asian region with vast diversities of socio-economic and networking conditions while inviting high-quality and recent research results from the global international research community.

After the success of the first AINTEC in 2005, the Organization Committee members agreed to continue AINTEC as annual events, and keep most of the committee members for the second year to foster the community. This year's conference was also jointly organized by the Internet Educational and Research Laboratory of the Asian Institute of Technology (AIT) and the WIDE Project with support from the APAN-TH community.

AINTEC 2006 solicited papers, among other things, on the survival of the Internet in order to provide alternative means of communication in emergency and chaotic situations in response to the recent natural disaster in Asia.

The main topics include: Mobile IP

Mobile Ad Hoc and Emergency Networks

Multimedia or Multi-Services IP-Based Networks

Peer-to-Peer

Measurement and Performance Analysis

Internet over Satellite Communications

These are the same topics as in AINTEC 2005, again to foster the community with focused research agendas. There were 36 submissions to the Technical Program, and we selected the 12 papers presented in these proceedings. In addition, we have five invited papers by leading experts in the field.

Finally, we would like to acknowledge the conference General Chair, Kanchana Kanchanasut of AIT, and the Local Organizers team from AIT, namely, Mohammad Abdul Awal, Withmone Tin Latt and Yasuo Tsuchimoto, for organizing and arranging this conference. We are also grateful to the French Ministry of Foreign Affairs through its French Regional Cooperation and the ICT Asia project (STIC-ASIA) for providing travel support.

November 2006

Kenjiro Cho
Philippe Jacquet
AINTEC 2006 Program Committee Co-chairs

Organization

General Chair

Kanchana Kanchanasut (Asian Institute of Technology, Thailand)

Program Committee Co-chairs

Kenjiro Cho, WIDE Project, Japan
Philippe Jacquet, INRIA, France

Program Committee

Alain Jean-Marie (LIRMM/INRIA, France)
Alexandru Petrescu (Motorola)
Anan Phonphoem (Kasetsart University, Thailand)
Antti Tuominen (Helsinki University of Technology, Finland)
Chalermek Intanagowiwat (Chulalongkorn University, Thailand)
Hyunchul Kim (CAIDA, USA)
Jun Takei (Intel KK, Japan)
Jochen Schiller (Freie Universität, Germany)
Kazunori Sugiura (Keio University, Japan)
Kenjiro Cho (IIJ, Japan)
Noel Crespi (INT, France)
Patcharee Basu (SOI/ASIA, Japan)
Poompat Saengudomlert (Asian Institute of Technology, Thailand)
Ryuji Wakikawa (Keio University, Japan)
Randy Bush (IIJ, USA)
Shigeya Suzuki (USC/ISI, USA)
Tapio Erke (Asian Institute of Technology, Thailand)
T.C. Wan (USM, Malaysia)
Teerapat Sa-nguankotchakorn (Asian Institute of Technology, Thailand)
Thierry Ernst (INRIA, France)
Thomas Clausen (Polytechnique, France)
Thomas Noel (University Louis Pasteur, France)
T.J. Kniveton (NOKIA Research Center, USA)
Yasuo Tsuchimoto (interELab, AIT, Thailand)
Yoshifumi Nishida (Sony CSL, Japan)
Youki Kadobayashi (NAIST, Japan)
Youngseok Lee (CNU, Korea)

Local Organization

Mohammad Abdul Awal (Asian Institute of Technology, Thailand)

Withmone Tin Latt (Asian Institute of Technology, Thailand)

Yasuo Tsuchimoto (Asian Institute of Technology, Thailand)

Table of Contents

Keynote Paper 1

An End-User-Responsive Sensor Network Architecture for Hazardous Weather Detection, Prediction and Response	1
<i>Jim Kurose, Eric Lyons, David McLaughlin, David Pepyne, Brenda Philips, David Westbrook, Michael Zink</i>	

Session 1: Service Architecture

On Scalability of DHT-DNS Hybrid Naming System.....	16
<i>Yusuke Doi, Shirou Wakayama, Masahiro Ishiyama, Satoshi Ozaki, Atsushi Inoue</i>	
A Topology-Aware Overlay Multicast Approach for Mobile Ad-Hoc Networks	31
<i>Mohamed Ali Kaafar, Cyrine Mrabet, Thierry Turletti</i>	
Next Generation Network Service Architecture in the IP Multimedia Subsystem	48
<i>Anahita Gouya, Noël Crespi, Lina Oueslati</i>	

Invited Paper 1

Live E! Project: Sensing the Earth	61
<i>Masaya Nakayama, Satoshi Matsuura, Hiroshi Esaki, Hideki Sunahara</i>	

Session 2: Multicast

Performance Analysis and Comparison Between Multicast and Unicast over Infrastructure Wireless LAN	75
<i>Anan Phonphoem, Suchaisri Li-On</i>	
NAT-PT with Multicast	90
<i>Maneenate Puongmanee, Robert Elz, Sinchai Kamolphiwong, Thossaporn Kamolphiwong, Chatchai Jantaraprim, Touchai Angchuan</i>	

Keynote Paper 2

Network Processing Hardware 103
Mary Inaba, Kei Hiraki

Session 3: Performance in WLAN

Channel Status Aware Proportional Fair Scheduling Algorithm for
IEEE 802.15.3 WPAN 113
Sung-Don Joo, Chae-Woo Lee

Application Performance Assessment on Wireless Ad Hoc Networks 128
*Răzvan Beuran, Ken-ichi Chinen, Khin Thida Latt,
Toshiyuki Miyachi, Junya Nakata, Lan Tien Nguyen,
Yoichi Shinoda, Yasuo Tan*

Proactive Resilience to Dropping Nodes in Mobile Ad Hoc Networks 139
Ignacy Gawędzki, Khaldoun Al Agha

Session 4: Routing

A New Hybrid Traffic Engineering Routing Algorithm for Bandwidth
Guaranteed Traffic 159
Zhaowei Meng, Jinshu Su, Stefano Avallone

Determining the Cause and Frequency of Routing Instability
with Anycast 172
James Hiebert, Peter Boothe, Randy Bush, Lucy Lynch

Invited Paper 3

Modeling Mobility with Behavioral Rules: The Case of Incident
and Emergency Situations 186
*Franck Legendre, Vincent Borrel, Marcelo Dias de Amorim,
Serge Fdida*

Session 5: Multihoming in Mobile Networks

Node Density Effects on Reactive Routing Protocols for Mobile
Ad Hoc Wireless Networks 206
*Hean-Loong Ong, Tat-Chee Wan, Hean-Kuan Ong,
Sureswaran Ramadass*

Building Fault Tolerant Networks Using a Multihomed Mobile Router: A Case Study	222
<i>Romain Kuntz, Jean Lorchat</i>	

Invited Paper 4

MANET and NEMO Converged Communication	235
<i>Kouji Okada, Ryuji Wakikawa, Jun Murai</i>	

Author Index	253
---------------------------	-----

An End-User-Responsive Sensor Network Architecture for Hazardous Weather Detection, Prediction and Response

Jim Kurose^{1,2}, Eric Lyons¹, David McLaughlin^{1,3}, David Pepyne¹, Brenda Philips¹,
David Westbrook¹, and Michael Zink^{1,2}

¹ Center for Collaborative Adaptive Sensing of the Atmosphere

² Department of Computer Science

³ Department of Electrical and Computer Engineering

University of Massachusetts

Amherst MA 10003 USA

{kurose@cs.umass.edu, elyons19@hotmail.com,

mclaughlin@ecs.umass.edu, pepyne@ecs.umass.edu,

bphilips@ecs.umass.edu, westy@cs.umass.edu, zink@cs.umass.edu}

Abstract. We present an architecture for a class of systems that perform distributed, collaborative, adaptive sensing (DCAS) of the atmosphere. Since the goal of these DCAS systems is to sense the atmosphere when and where the user needs are greatest, end-users naturally play the central role in determining how system resources (sensor targeting, computation, communication) are deployed. We describe the meteorological command and control components that lie at the heart of our testbed DCAS system, and provide timing measurements of component execution times. We then present a utility-based framework that determines how multiple end-user preferences are combined with policy considerations into utility functions that are used to allocate system resources in a manner that dynamically optimizes overall system performance. We also discuss open challenges in the networking and control of such end-user-driven systems.

Keywords: Sensor networks, collaborative adaptive atmospheric sensing, end-user utility.

1 Introduction

Over the past twenty years, the use of networked computing systems has evolved from being primarily general/multi-purpose in nature to often being highly specialized and highly mission-specific. Perhaps nowhere is this trend clearer than in the case of sensor networks [DDAS 2006, Chong 2003, Estrin 2002] – a combined sensing, communication, and computing infrastructure designed to measure, monitor, predict, and (in some cases) control a particular environment or process. Such mission-specific computing naturally implies an increased emphasis on the end-user – the *raison d’être* for the system in the first place – and on tailoring and optimizing system operation to meet end-user-defined, application-specific goals. Although it is thus

important to engineer sensor network systems to meet these specific end-user needs and application scenarios, it is also important to avoid a highly-specialized “stovepipe” design that is difficult to evolve for use in new (or changed) application scenarios or end-user/application requirements.

In this paper, we describe the software architecture for a sensor network system consisting of a relatively small number (tens) of low-power X-band radars that detect and predict hazardous weather via distributed, collaborative, adaptive sensing (DCAS) of the lowest few kilometers of the earth’s atmosphere. The architecture finds use in a variety of application scenarios including hazardous wind (e.g., tornado) and precipitation sensing, in both resource-rich and resource-challenged environments. *Distributed* refers to the use of a number of small radars, spaced close enough to “see” close to the ground in spite of the Earth’s curvature and avoid resolution degradation caused by radar beam spreading. *Collaborative* operation refers to the coordination (when advantageous) of the beams from multiple radars to view the same region in the atmosphere, thus achieving greater sensitivity, precision, and resolution than possible with a single radar. *Adaptive* refers to the ability of these radars and their associated computing and communications infrastructure to dynamically reconfigure in response to changing weather conditions and end-user needs. As part of our activities in the NSF Engineering Research Center for Collaborative Adaptive Sensing of the Atmosphere (CASA) [McLaughlin 2005, CASA 2006a], we have instantiated the DCAS paradigm in an operational testbed for hazardous wind sensing in southwestern Oklahoma, and have several additional DCAS systems under development [Donovan 2005b].

The goal of these DCAS systems is to sense the atmosphere *when and where the user needs are greatest*. Thus, user requirements naturally play *the* central role in determining how system resources (sensor targeting, computation, communication) are deployed at any given point in time. Our DCAS systems adopt a utility-based framework in which multiple end-user preferences are combined with policy considerations into utility functions that are used to allocate system resources in a manner that dynamically optimizes overall system performance. In this paper, we described the software architecture of our DCAS systems, focusing on the driving role of end-user considerations in system design and operation. We also discuss open challenges in the networking and control of such end-user-driven systems.

The remainder of this paper is structured as follows. In section 2, we overview the DCAS software architecture, highlighting the crucial role played by the end users. In section 3, we discuss how end-user preferences and policy considerations are mapped to the utility functions that then determine how system resources are allocated. We discuss our experiences with initial end-user evaluations, and how these experiences have led to changes in our utility functions. In section 4 we discuss several open challenges in the networking area, focusing on those that arise from end-user considerations. Section 5 concludes this paper.

2 Overview of DCAS Software Architecture

Figure 1 shows the overall software architecture of the meteorological command and control (MC&C) components that lie at the heart (or perhaps more appropriately, the

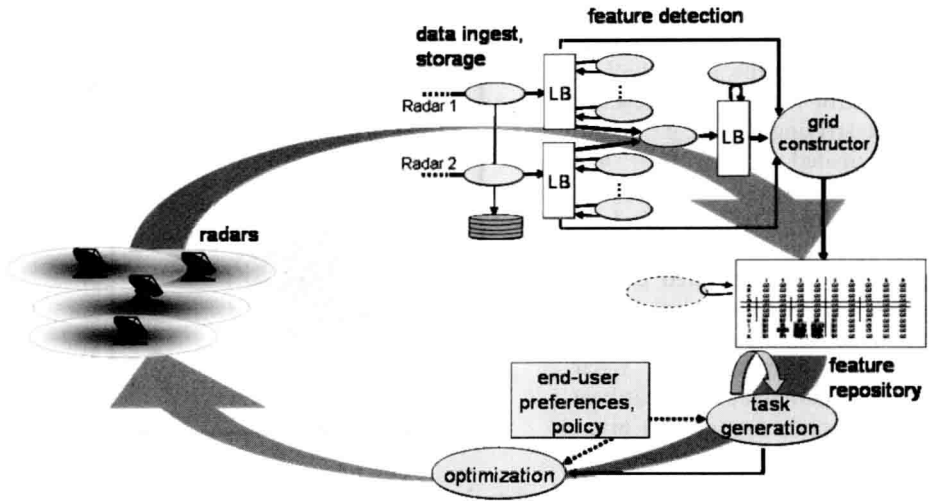


Fig. 1. Software architecture for meteorological command and control (MC&C)

“brains”) of our DCAS system. The operational version of the MC&C is currently centralized, but we are currently implementing a distributed version of the MC&C as well. The main system control loop divides into loosely-coupled upper and lower halves. In the upper loop, data is ingested from the remote radars, meteorological features are identified in the data, and higher-level meteorological features (e.g., wind rotation, wind shear and rotational divergence, areas of high precipitation) are posted on a feature repository, from which they (and the underlying radar data) are made available to end users. The shaded ovals in the feature-detection portion of the MC&C are the individual detection algorithms that detect specific features in the data. These modules are connected together by an event-based broadcast publication/subscribe mechanism known as a linear buffer [Hondl 2003] that notifies downstream modules when data or events are available for processing. The feature repository is a blackboard-like module [Jaganathan 1989] that allows data to be asynchronously written (posted) and read. The lower half of the control loop uses the posted features and end-user preferences and policy to identify areas of meteorological interest in the radars’ footprints (task generation) and then optimizes the configuration of each radar (i.e., the location and width of the sector within its footprint to be scanned), as discussed in detail in Section 3.

The system operates on a 30-second “heartbeat,” with the upper control loop in Figure 1 (data ingest and feature detection) proceeding *asynchronously* from the lower control loop (generation of radar commands for the next 30-second heartbeat). During each heartbeat, data is ingested and processed in streaming mode. Δ seconds before the end of each heartbeat (a value of Δ equal to approximately 4 seconds is currently used in our testbed), command generation components begin their computation to determine how to best target the radars for the next 30-second heartbeat based on data currently posted in the feature repository. This temporal

decoupling of the upper and lower halves of the control loop avoids stalling the lower half of the control loop in the presence of late-arriving data (e.g., due to unanticipated network and processing delays elsewhere in the system). Thus, while late-arriving data may affect the quality of the computed control (since the targeting decision is based on older data), it will not delay the generation of radar targeting commands. A heartbeat of 30 seconds was chosen given the coverage of each radar (a radius of approximately 30km), the timescale at which meteorological features evolve, and the operating characteristics of the mechanically-scanned radars currently in our testbed. The value of Δ will clearly play an important role in how well the radar network is able to sense the environment. The smaller the value of Δ , the more recent the data that will be used to re-target the radars. We report on the measured runtimes of various system components below that indicate that a value of Δ of approximately 4 seconds is appropriate in our operational system.

We have completed an initial four-radar testbed that covers a 7,000 square km region in southwestern Oklahoma, a region that receives an average of four tornado warnings and 53 thunderstorm warnings per year. Figure 2 shows the radar locations, including their dual Doppler (overlap) regions. Data is streamed from each of the four radars at up to 1 Mbps, over a pre-provisioned network with adequate capacity to handle this peak rate, to a central site where the MC&C operates. (Thus, in this resource-rich initial testbed, network concerns such as routing and congestion control are not of significant concern; we discuss networking challenges in resource-poor DCAS system in Section 4). The MC&C itself executes on a 3-node (3.2 GHz Intel CPU, 1 GB RAM) cluster with 4 TB of storage; the cluster can be easily scaled by adding additional processors if additional computing or storage resources are needed. The end-users of this testbed are the National Weather Service (NWS) Forecast Office in Norman, OK, a group of emergency managers who have jurisdictional authority within and upstream of the testbed area, several private sector entities, and CASA’s science researchers themselves.

Given the soft-real-time nature of this system, it is crucial that the execution times of the various software components shown in Figure 1 fit well into the fundamental 30-second heartbeat. Figure 3 shows the execution times of selected MC&C components: per-elevation execution times for data ingest, a peak echo detection algorithm, and task generation and optimization. These execution times were measured in our 4-node testbed during a period of time when several storm cells were moving through the testbed’s coverage area. The data ingest algorithm receives data in an elevation scan

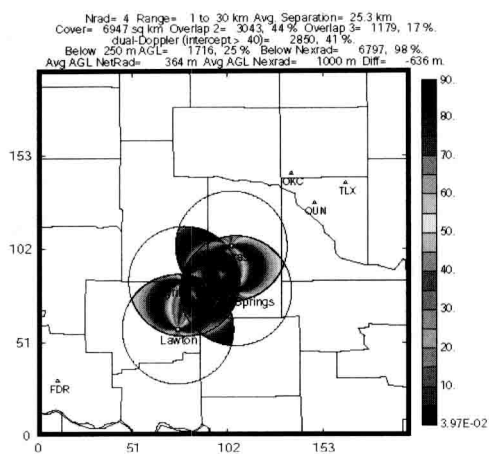


Fig. 2. Radar siting in southwestern OK, showing dual-Doppler regions

from a radar (there may be several sequential elevation scans within one 30-second heartbeat), performs simple cleanup (thresholding) of the data, sequentially stores the reflectivity, wind-velocity and other data in separate files in a common format, and notifies the downstream detection modules of the availability of data via the broadcast pub/sub mechanism. The peak echo algorithm identifies the region of peak reflectivity within a radar's footprint. Since the data ingest algorithm operates in a streaming manner, and since peak echo detection, task generation, and optimization components all execute in subsecond times, these execution times are relatively small with respect to the 30-second heartbeat interval. Execution times from other MC&C components (using emulated radar input data) have been report in [Zink 2005].

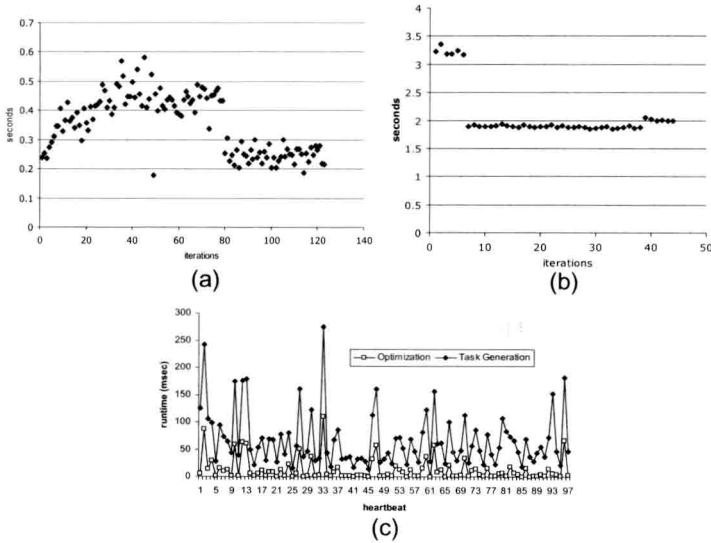


Fig. 3. Execution runtimes for (a) data ingest, (b) peak echo detection, and (c) task generation/optimization

3 Responsiveness to End User Requirements

Our discussion above has focussed on the overall DCAS software system architecture and software component runtimes within the 30-second heartbeat. We now turn our attention to *the* central aspect of DCAS control – determining *where* and *how* the radars should target their scans during a 30-second interval. This is the role of the task generation and optimization components of the MC&C and it is here that end-user considerations are of central concern.

3.1 The MC&C Equation

In order to simplify our discussion below, we consider the “simpler” problem of determining only *where* the radars should focus their scanning, ignoring additional

considerations such as the most appropriate radar waveform to be employed, the pulse repetition frequency (PRF), and more. This simple setting will nonetheless allow us to illustrate the key end-user-related aspects of the MC&C.

We define a *radar configuration* to be the set of sectors to be scanned by the radars during a 30-second interval. The MC&C takes a *utility-based approach* towards determining the optimal radar configuration for a particular interval, where (as we will see) utility is directly defined in terms of end-user requirements. Two sets of factors contribute to the utility of a particular radar configuration. The first set of factors is concerned with *how well* a particular portion of the atmosphere is sensed by a given radar configuration. The second set of factors is concerned with *how important* the scanned sectors are to the end users. The overall utility of a particular configuration combines these two considerations of “how well” and “how important.”

Before specifying how utility is computed, let us first qualitatively discuss the factors that will come into consideration. The first set of factors is related to the quality of the data from the scanned sectors (“how well”):

- **Scan sector size.** Each radar can be tasked to scan a sector with an arbitrary starting point and at any width from 60 degrees to 360 degrees during the 30-second heartbeat. Scanning a smaller sector allows more sensing energy to be focussed into that sector, either by dwelling on each radial position for a longer period of time (than in the case of a wider scan) and hence obtaining more accurate estimates of sensed values [Donovan 2005a], or by scanning a larger number of elevations. The increased “quality” of sensed data within a narrowly-scanned sector must then be weighed against obtaining no data from those portions of the radar’s coverage that were not scanned during the heartbeat.
- **Coordination among radars.** It is often advantageous to have two or more radars focus their scans on overlapping regions in the atmosphere, a so-called dual-Doppler region. One radar may be able to “see” that portion of the atmosphere better (e.g., due to less signal attenuation), and data from multiple radars allows for more accurate estimation of wind velocity vectors. Another benefit of radar coordination arises even when each meteorological feature is scanned by a single radar – when a particularly high utility feature can be scanned by more than one radar, the system can scan that feature with the radar that allows remaining radars to scan other meteorological features of next highest utility.

The second set of factors is concerned with *how important* the scanned portions of the atmosphere are to the end-users:

- **Expressing the preferences of multiple end users.** The initial end users in our Oklahoma testbed are (i) the National Weather Service (NWS) forecasters, both in the Norman, Oklahoma Forecast Office and in the NWS Warning Decision Training Branch which trains forecasters nationally, (ii) a group of emergency managers who have jurisdictional authority within and upstream of the test bed area, and (iii) CASA’s science researchers themselves. Different end users will derive different utility from a particular radar configuration. For example, an NWS forecast office may use 360 degree sweeps of DCAS radar data close to the ground to monitor the evolution of a storm to determine whether to issue a

tornado warning, while at the same time an emergency manager may require two collaborating radars to locate precisely the most intense part of a storm for public notification or for deploying weather spotters. A researcher initializing a numerical prediction model needs 360 degree scans at all elevations. Each of the end user groups must be able to express their relative preferences among different radar configurations, and these preferences must then be incorporated into system operation.

- **Policy: mediating among the conflicting scanning requests of different end users.** The differing information requirements of different end users means that it may not be possible to satisfy the needs of all end users at the same time. If a given radar configuration is of particularly high utility to one group of users but of lower utility to another group of users, while a second configuration is of low utility to the first group of users but of high utility to the second group of users, the system must decide which of these two alternatives is preferable. That is, a *policy* mechanism must be defined for mediating among the conflicting scanning demands of different end-users.

The above considerations give rise to the following optimization problem (which we refer to as the MC&C equation) that must be solved by the MC&C:

$$J = \max_{\text{configurations}, C} \sum_{\text{tasks}, t} U(t, k) Q(t, C) \quad (1)$$

where:

- t is the set of so-called *tasks* – meteorological features that are within the radars’ coverage areas and thus can potentially be scanned at time interval k . Examples of tasks include areas of wind rotation, areas of high reflectivity, and areas of wind shear. These tasks are created by the task-generation module shown in Figure 1, based on detected features posted in the feature repository. As we saw in Figure 3, the set of tasks in our testbed can be generated in less than a second by the task-generation module.
- $U(t, k)$ is the *aggregate end-user utility* of task t to the set of end users, and captures the “how important” aspect of utility. $U(t, k)$ in turn is defined as:

$$U(t, k) = \sum_{\text{groups}, g} w_g U_g(t, k)$$

where g is the set of user groups, $U_g(t, k)$ is the utility of task t at time interval k to user group g , and w_g is the weight associated with user group g . In section 3.2 below, we discuss how $U_g(t, k)$ is computed. w_g is a value between 0 and 1, reflecting the relative priority of user group g with respect to other user groups. The values of w_g are set by policy. Note that the architecture itself is policy-neutral in that it does not prescribe values for w_g but instead provides a *mechanism* for weighting the relative importance of different user groups.

- $Q(t, C)$ is the *quality* of the scanning of task t under radar configuration C , capturing the “how well” aspect of utility. $Q(t, C)$ is a value between 0 and 1, with a value of 1 representing the highest quality possible. The details of the computation of $Q(t, C)$ can be found in [Pepyne 2006]. We note here that for a particular task, t , the value of $Q(t, C)$ (i) increases as the physical location

(center) of the task becomes closer to a radar (ii) increases as increasingly more of the physical extent of the task is covered by a radar in configuration C , and (iii) increases as function of the number of radars whose scans cover the task.

Several comments are in order regarding the MC&C equation. First, note that the form of the MC&C equation involves a sum over all tasks, where the utility of each task has a quality component, $Q(t, C)$ (reflecting “how well” the task is sensed by the radars in a particular configuration) and an end-user utility component, $U(t, k)$ (reflecting how important that task is to the end users). The overall utility of a given task is the product of these two components. Rather than taking the product of these two components, an alternate approach would be to allow end-user utility itself to depend on the quality component, i.e., to define task utility under configuration C in the form $U(t, k, Q(t, C))$, rather than $U(t, k)Q(t, C)$ as in Equation 1. This alternative formulation would have resulted in a significantly more complex optimization problem, since end-user utility would be a function of the scan quality of the configuration. Our separation of “how well” and “how important” into two independent considerations represents an architectural decision to separate lower-level, radar-specific, sensing considerations from higher-level end-user considerations. This also illustrates one way in which we have avoided building a stovepipe system. If a new set of radars, with new operating characteristics were to be used, we need only change $Q(t, C)$. On the other hand, with a task utility of the form $U(t, k, Q(t, C))$, we would have to redefine end-user utility as well – a difficult and time-consuming task that requires numerous interactions with the end users, as we will see in Section 3.2. We emphasize, however, that both task quality *and* task importance are taken into account in the end-user utility calculation (through their product).

Second, we note that while it appears that the optimization is considering each 30-second heartbeat as an independent optimization problem, this is actually not the case. As we will see in section 3.2, users have specified not only *what* tasks they want scanned, but also how *often* different types of tasks need to be scanned. As we will see, a typical end-user rule has the form: “*If a meteorological phenomena of type X is detected, then scan it with Y radars (when possible) at least once every Z heartbeats.*” To implement this capability we keep track of the *time-since-last-scanned* for each task. If the *time-since-last-scanned* is less than the user-defined interval between scans (Z heartbeats in the previous sentence), then the task has a base utility value. If the *time-since-last-scanned* is greater than the user-defined interval, we scale (increase) the utility value $U_g(t, k)$ of task t at each heartbeat until the task is eventually scanned. At this point, the task’s utility value is set to its base value, and the scaling process begins again. With this simple scheme, we implicitly solve what is otherwise a complex multistage optimization problem. [Manfredi 2006] examines a non-myopic multistage optimization formulation of the MC&C optimization problem.

3.2 End-User Utility

Given the importance of end-users in our DCAS systems, a crucial challenge is to define the end-user utility functions, $U_g(t, k)$, in a manner that is consistent with the