STROBE MVS

STROBE CA-IDMS Feature

Release 3.0

Please direct questions about STROBE MVS
or comments on this document to:

STROBE MVS Technical Support
Compuware Corporation
124 Mount Auburn Street
Cambridge MA 02138-5758
1-800-585-2802 or
1-617-661-3020
1-617-498-4010 (fax)
strobe-sup@compuware.com

Outside the USA and Canada, please contact
your local Compuware office or agent.

This document and the product referenced in it are subject to the following legends:

Copyright 2002 Compuware Corporation. All rights reserved. Unpublished rights reserved under the Copyright Laws of the United States.

U.S. GOVERNMENT RIGHTS-Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in Compuware Corporation license agreement and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (OCT 1988), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable. Compuware Corporation.

This product contains confidential information and trade secrets of Compuware Corporation. Use, disclosure, or reproduction is prohibited without the prior express written permission of Compuware Corporation. Access is limited to authorized users. Use of this product is subject to the terms and conditions of the user's License Agreement with Compuware Corporation.

STROBE, iSTROBE and APMpower are trademarks or registered trademarks of Compuware Corporation.

The following are trademarks of Computer Associates International, Inc: CA-eMail+ CA-MIM, and CA-Multi-image Manager. The following are registered trademarks of Computer Associates International, Inc: CA-ACF2, CA-IDMS, CA-Optimizer, CA-OptimizerII, CA-Panexec, and CA-Top Secret. Acrobat® Reader copyright © 1987-1998 Adobe Systems Incorporated. All rights reserved. Adobe, Acrobat, and Acrobat Reader are trademarks of Adobe Systems Incorporated. AD/Cycle, BookManager, CICS, DB2, IBM, IMS/ESA, Language Environment, MQSeries, OS/2, Software Mall, VisualGen, and VTAM are trademarks of International Business Machines Corporation. Microsoft is a registered trademark of Microsoft Corporation. Windows, Windows NT, and Windows 98 are trademarks of Microsoft Corporation. Attachmate and EXTRA! are registered trademarks of Attachmate Corporation. RUMBA is a registered trademark of Wall Data Inc. WRQ and Reflection are registered trademarks of WRQ, Inc.

Adobe ® Acrobat ® Reader copyright © 1987-2001 Adobe Systems Incorporated. All rights reserved. Adobe and Acrobat are trademarks of Adobe Systems Incorporated.

All other company and product names are trademarks or registered trademarks of their respective owners.

Doc. CWSTUR3A
January 21, 2003.

Contents

Figures v

Summary of Changes vii
Changes to the STROBE CA-IDMS Feature vii
Changes to this Manual vii

Introduction ix
How This Manual Is Organized ix
How to Use This Manual ix
The STROBE Library ix
STROBE Feature Manuals x
Online Documentation x
Online Help x
Other Compuware Application Performance Management Products xi
iSTROBE xi
SQL Analysis Feature xi
APMpower xi
Compuware APM Technical Support xi
Compuware APM Training xi
Compuware APM Service Offerings xii
APM Installation Assurance xii
Application Performance Management Consulting xii
Application Performance Assessment xii

Chapter 1. Overview 1-1
CA-IDMS Capabilities 1-1
CA-IDMS Environments 1-2
CA-IDMS-DC/CV Mode 1-2
CA-IDMS Central Version Batch Mode 1-3
CA-IDMS Local Mode Batch 1-4
CA-IDMS-DC/CV-UCF Mode 1-4
Identification of Activity in Load Modules 1-5
Identification of Activity in Taskcodes 1-5
CA-IDMS Pseudo-Taskcodes (Transactions) 1-5
STROBE-Generated Module Names 1-5
Other Pseudo-Names 1-5
Benefits of the STROBE CA-IDMS Feature 1-6

Chapter 2. Using the STROBE CA-IDMS Feature 2-1
Invoking the STROBE CA-IDMS Feature 2-1
CA-IDMS-DC/CV Online Regions 2-1
CA-IDMS Central Version and Local Mode Batch Regions 2-2
Indexing ADS/O Dialogs 2-2
Producing ADS/O Dialog Listing Data Sets 2-3
Using STROBE/ISPF to Index ADS/O Dialogs 2-4
Using the STROX Procedure to Index ADS/O Dialogs 2-4
Using Your Own Job Control Statements 2-5
Tailoring STROBE Reports 2-5
Suppressing Attribution Reports in the Performance Profile 2-5
Suppressing Attribution Reports with Low Activity 2-5
Suppressing All Attribution Reports 2-6

Suppressing CA-IDMS-Specific Attribution Reports 2-6

Chapter 3. The Performance Profile for Online Regions 3-1
Measurement Session Data Report 3-1
Choosing Between the Execution and Wait Reports 3-1
Reports That Show CPU Execution 3-2
Program Module Reports 3-2
ADS/O Dialog Reports 3-2
Attribution Reports 3-2
Transaction Reports 3-2
Program Module Reports 3-3
Program Section Usage Summary Report 3-3
Program Usage by Procedure 3-4
Activity in System Modules 3-4
Activity in Application Programs 3-4
Dialog Reports 3-5
Dialog CPU Summary Report 3-5
CPU Usage by ADS/O Statement 3-6
Dialog Identification 3-7
Statement Number 3-7
Statement Text 3-8
Attribution Reports 3-9
Attribution of CPU Execution Time 3-9
Header Lines 3-9
Detail Lines 3-9
Total Line 3-10
Transaction Reports 3-11
Transaction Summary 3-11
Transaction Usage by Control Section Reports 3-12
Wait Time by Module Report 3-15
Reports That Show I/O Utilization 3-16
Resource Demand Distribution 3-16
I/O Facility Utilization 3-17
Data Set Characteristics 3-17

Chapter 4. The STROBE Performance Profile for Batch Regions 4-1
Measurement Session Data Report 4-1
Wait Time by Module 4-2
Attribution of CPU Wait Time Report 4-2
Attribution of CPU Execution Time Report 4-4
Using Database Analysis Utilities with STROBE 4-4
Sorted Sets 4-4
Indexed Sets 4-5
CALC or VIA 4-5
Database Buffering 4-5

Index I-1

Figures

1-1. MVS Environment Running Jobs that use CA-IDMS 1-2
2-1. STROBE - DATA COLLECTORS Panel....2-2
2-2. STROBE - INDEX TO CREATE A MAP DATA SET Panel....2-4
2-3. STROBE - TAILOR REPORTS Panel 2-6
3-1. Measurement Session Data Report 3-1
3-2. Program Section Usage Summary Report....3-3
3-3. Program Usage by Procedure Report fo CA-IDMS Modules....3-5
3-4. Dialog CPU Summary Report 3-6
3-5. CPU Usage by ADS/O Statement Report 3-7
3-6. ADS/O Dialog Code Showing CONCAT....3-8
3-7. ADS/O Dialog Code Showing ACCEPTS....3-8
3-8. Attribution of CPU Execution Time Report 3-10
3-9. Transaction Summary Report....3-12
3-10. Transaction Usage by Control Section Report for ADS/O Tasks....3-13
3-11. Transaction Usage by Control Section Report for CA-IDMS Application Tasks....3-14
3-12. Transaction Usage by Control Section Report for Central Version Batch Jobs 3-14
3-13. Transaction Usage by Control Section Report for CA-IDMS Tasks....3-15
3-14. Wait Time by Module Report 3-16
3-15. Resource Demand Distribution Report 3-16
3-16. I/O Facility Utilization Summary Report 3-17
3-17. Data Set Characteristics Report 3-18
4-1. Measurement Session Data Report 4-1
4-2. Wait Time by Module Report for Batch Regions....4-2
4-3. Attribution of CPU Wait Time Report....4-3
4-4. CA-IDMS/COBOL Code Showing the DML Sequence Numbers 4-3
4-5. Attribution of CPU Execution Time Report 4-4

Summary of Changes

This section discusses the changes to the STROBE CA-IDMS Feature for STROBE MVS for Sysplex Release 3.0.

Changes to the STROBE CA-IDMS Feature

The STROBE CA-IDMS Feature did not change from STROBE MVS for Sysplex Release 2.5.0.

Changes to this Manual

The technical content of this manual did not change for STROBE MVS for Sysplex Release 3.0.

Introduction

This manual describes measurement concepts applicable to and specific data made available by the STROBE CA-IDMS Feature of the STROBE MVS Application Performance Measurement System. The STROBE CA-IDMS Feature augments functions provided by the basic STROBE system.

The STROBE MVS Application Performance Measurement System and the STROBE CA-IDMS Feature are products designed for IBM MVS/ESA, IBM OS/390, and IBM z/OS systems. The STROBE CA-IDMS Feature is designed for use with releases 10.2, 12.0, 12.01, 14.0 and 15.0 of Computer Associates International, Inc.'s Integrated Database Management System Central Version (CA-IDMS-CV), the CA-IDMS-DC Teleprocessing Monitor, and Application Development System/Online (ADS/O).

How This Manual Is Organized

Chapter 1, “Overview” is an introduction to the STROBE CA-IDMS Feature.

Chapter 2, “Using the STROBE CA-IDMS Feature” describes how to use the STROBE CA-IDMS Feature to measure CA-IDMS regions.

Chapter 3, “The Performance Profile for Online Regions” explains how to interpret the STROBE Performance Profile reports that show CPU usage data for CA-IDMS online regions.

Chapter 4, “The STROBE Performance Profile for Batch Regions” explains how to interpret the Performance Profile for CA-IDMS batch regions.

How to Use This Manual

You should read Chapter 1, “Overview” and Chapter 2, “Using the STROBE CA-IDMS Feature” before submitting a measurement request. To interpret a Performance Profile for an online CA-IDMS region, read Chapter 3, “The Performance Profile for Online Regions”. To interpret a Performance Profile for a batch CA-IDMS region, read Chapter 4, “The STROBE Performance Profile for Batch Regions”.

The STROBE Library

The STROBE base product manuals include:

- *STROBE MVS Concepts and Facilities*, document number CWSTGX3A

  STROBE MVS Concepts and Facilities explains how to decide which programs and online regions to measure, when to measure them, and how to interpret the reports in the STROBE Performance Profile.

- *STROBE MVS Messages,* document number CWSTXM3A

  STROBE MVS Messages lists all messages and abnormal termination (ABEND) codes, describes how to interpret them, and in many cases suggests a corrective action.

- *STROBE MVS System Programmer’s Guide*, document number CWSTXI3A

The *STROBE MVS System Programmer's Guide* explains how to install and maintain STROBE.

- *STROBE MVS User's Guide*, document number CWSTUX3A and the *STROBE MVS User's Guide with Advanced Session Management*, document number CWSTUA3A

  The *STROBE MVS User's Guide* explains how to use STROBE to measure application performance. The *STROBE MVS User's Guide with Advanced Session Management* explains how to use STROBE with the STROBE Advanced Session Management Feature to measure application performance. Users who have the STROBE Advanced Session Management Feature will use this manual rather than the *STROBE MVS User's Guide*.

- *STROBE MVS Application Performance Measurement System Quick Reference*

  The *STROBE MVS Application Performance Measurement System Quick Reference* is a convenient reference for how to use STROBE and for interpreting the STROBE Performance Profile.

STROBE Feature Manuals

These manuals describe the optional features of the STROBE MVS Application Performance Measurement System. Each manual describes measurement concepts applicable to and specific data made available by the feature.

- *STROBE MVS User's Guide with Advanced Session Management*, document number CWSTUA3A
- *STROBE ADABAS/NATURAL Feature*, document number CWSTUN3A
- *STROBE CA-IDMS Feature*, document number CWSTUR3A
- *STROBE CICS Feature*, document number CWSTUC3A
- *STROBE COOL:Gen Feature*, document number CWSTUG3A
- *STROBE CSP Feature*, document number CWSTUP3A
- *STROBE DB2 Feature,* document number CWSTUD3A
- *STROBE IMS Feature,* document number CWSTUI3A
- *STROBE Interface Feature*, document number CWSTUF3A
- *STROBE Java Feature,* document number CWSTUJ3A
- *STROBE MQSeries Feature,* document number CWSTUM3A
- *STROBE UNIX System Services Feature,* document number CWSTUU3A

Online Documentation

STROBE manuals are available in HTML, Adobe Acrobat PDF format, and IBM BookManager format, on CD-ROM and at Compuware's technical support Web site at http://frontline.compuware.com.

Online Help

STROBE products provide the following online information:

- STROBE/ISPF Online Tutorials, Option T from the STROBE/ISPF STROBE OPTIONS menu
- STROBE/ISPF Online Message Facility, Option M from the STROBE/ISPF STROBE OPTIONS menu

Other Compuware Application Performance Management Products

The following products and features work in conjunction with the STROBE MVS Application Performance Measurement System. These tools extend the benefits of application performance management (APM).

iSTROBE

iSTROBE enables you to view and analyze STROBE Performance Profile data on a workstation using a standard Web browser. Easy to install and easy to use, iSTROBE guides you through the performance analysis process and offers recommendations for improving performance. iSTROBE simplifies the performance analysis of applications that you measure with STROBE. For more information on iSTROBE, see the *iSTROBE Getting Started Guide*.

SQL Analysis Feature

The SQL Analysis Feature works in conjunction with STROBE and iSTROBE or APMpower to supply access path analyses and database and SQL coding recommendations for DB2 applications measured by STROBE. The SQL Analysis Feature pinpoints the most resource-consumptive static or dynamic SQL statements, explains why these statements might be inefficient, and provides recommendations to improve the performance of the DB2 application. For more information on the SQL Analysis Feature, see the *STROBE MVS User's Guide* or the *STROBE MVS User's Guide with Advanced Session Management*.

APMpower

The APMpower Application Performance Analysis System extends the benefits of STROBE to application developers who use workstations to develop, test, and maintain MVS applications. Developers employ the APMpower graphical user interface and advanced analytical aids to navigate the Performance Profile, analyze and improve application performance, and share performance knowledge across the IS organization. For more information about APMpower, see the APMpower documentation.

Compuware APM Technical Support

For North American customers, for technical support, please contact the Technical Support department by telephone at (800) 585-2802 or (617) 661-3020, by fax at (617) 498-4010, or by e-mail at strobe-sup@compuware.com.

To access online technical support, visit Compuware's FrontLine page on the World Wide Web at http://frontline.compuware.com and select the product "STROBE and APMpower."

For other international customers, please contact your local Compuware office or STROBE supplier.

Compuware APM Training

Compuware's Education Resources Group offers a range of training options for organizations that use STROBE, iSTROBE, and APMpower. To arrange Application Performance Management training, please contact Compuware at 1-800-835-3190 or visit Compuware's Education Resources Group at http://www.compuware.com/training

For other international customers, please contact your local Compuware office or STROBE supplier for a complete list of APM Training offerings.

Compuware APM Service Offerings

For North American customers, for information about current service offerings, please contact your local Compuware sales office or call Compuware Corporate Headquarters at 1-800-COMPUWARE (266-7892) or visit Compuware's APM Product page on the World Wide Web at http://www.compuware.com/products/strobe.

For other international customers, please contact your local Compuware office or STROBE supplier for a complete list of Services offerings.

APM Installation Assurance

The APM Installation Assurance service assists you in planning for, installing, customizing and using APM products. The service will help you maximize the value and benefits derived from the APM product family.

Consulting engineers work closely with your IT personnel to understand your operating environment and your organization's APM goals. The engineer will assist you in developing a customization and installation plan for STROBE, iSTROBE, and APMpower. The engineer will oversee the installation process and verify product readiness. The engineer will also help set up measurement request schedules, request groups, history records, AutoSTROBE measurement requests, and will verify the installation of the SQL Analysis Feature.

With APM Installation Assurance services, your organization can immediately maximize the value received from your investment in the APM product family. You will also benefit from a fully customized installation that will enhance the product functionality and increase the automation aspects of your APM initiatives.

Application Performance Management Consulting

The Application Performance Management (APM) Consulting services assist you in identifying and resolving specific performance problems in your OS/390 business-critical applications.

Using STROBE, iSTROBE, and APMpower, consulting engineers work closely with your IT personnel to measure an application's performance, identify performance improvement opportunities and make recommendations for implementing solutions.

With APM Consulting services, your organization cannot only resolve problems quickly and effectively, but also gain the skills necessary to prevent application performance degradation in the future.

Application Performance Assessment

The Application Performance Assessment (APA) service assists you in achieving a higher level of performance for your OS/390 business-critical applications.

Using STROBE, iSTROBE, and APMpower, consulting engineers work closely with your IT personnel to evaluate the efficiency of business-critical applications, identify opportunities for improving performance and document the potential savings that can result from implementing recommended solutions.

With APA services, you cannot only improve application performance quickly and effectively, but also gain the knowledge and skills necessary to implement and sustain a process-oriented application performance management (APM) program.

Chapter 1. Overview

The STROBE CA-IDMS Feature is a product that extends the basic STROBE functionality. This manual describes the various uses of the STROBE CA-IDMS Feature and how it works with STROBE to measure CA-IDMS applications.

CA-IDMS Capabilities

With the STROBE CA-IDMS Feature, you can measure CA-IDMS online and batch regions and evaluate the performance of

- ADS/O dialogs
- online transactions and external run units
- CA-IDMS online application programs written in assembler, COBOL, and PL/I
- CA-IDMS subsystem nucleus programs

You can also measure CA-IDMS batch application programs written in assembler, COBOL, and PL/I, which execute in either Central Version or local mode.

The STROBE CA-IDMS Feature dynamically obtains measurement information to produce the STROBE Performance Profile, which includes reports that

- show CPU usage by load module and control section
- identify CPU activity of CA-IDMS system and application tasks (transactions)
- attribute CPU time to ADS/O dialogs, processes, and commands
- attribute CPU time to online programs and external run units
- attribute CPU time in CA-IDMS batch database application programs to database DML and SQL commands
- provide function descriptors for CA-IDMS modules

The STROBE CA-IDMS Feature lets you identify the CA-IDMS application and system tasks that most affect the performance of your CA-IDMS system. Additionally, you can determine which CA-IDMS nucleus programs execute on behalf of those tasks. Using this knowledge, you can make decisions about how to

- streamline application code and request ADS/O and CA-IDMS services more efficiently
- concentrate efforts to reduce batch database application run time
- divide applications among CA-IDMS regions
- distribute hardware resources more efficiently for CA-IDMS applications

The STROBE CA-IDMS Feature identifies tasks that perform CA-IDMS system services by assigning them transaction names. It assigns these names based upon the

- CA-IDMS-DC taskcode
- external batch program (run-unit) name

The STROBE CA-IDMS Feature identifies CA-IDMS system tasks by assigning them pseudo-transaction names that reflect the types of services that the tasks provide. A complete list of these pseudo-transaction names appears in the section titled “CA-IDMS Pseudo-Taskcodes (Transactions)” on page 1-5.

This feature also identifies CA-IDMS nucleus and application programs by load module name, when available, and by a pseudo-load module name when not available.

CA-IDMS Environments

CA-IDMS applications execute in four different environments, each geared to a specific type of application processing need. The environments are:

- CA-IDMS-DC/CV mode
- CA-IDMS Batch Central Version mode
- CA-IDMS local mode batch
- CA-IDMS-DC/CV-UCF (Universal Communications Facility) mode

Figure 1-1 illustrates a possible job mix showing the four CA-IDMS environments. The organization of components within each region represents the high level interaction among them.

CA-IDMS-DC/CV Mode

CA-IDMS-DC/CV is an online region consisting of a teleprocessing component, CA-IDMS-DC, and a database management system, CA-IDMS Central Version (CV). CA-IDMS-DC provides a set of services that are requested by applications through DC Data Manipulation Language (DML) commands.

Figure 1-1. MVS Environment Running Jobs that use CA-IDMS

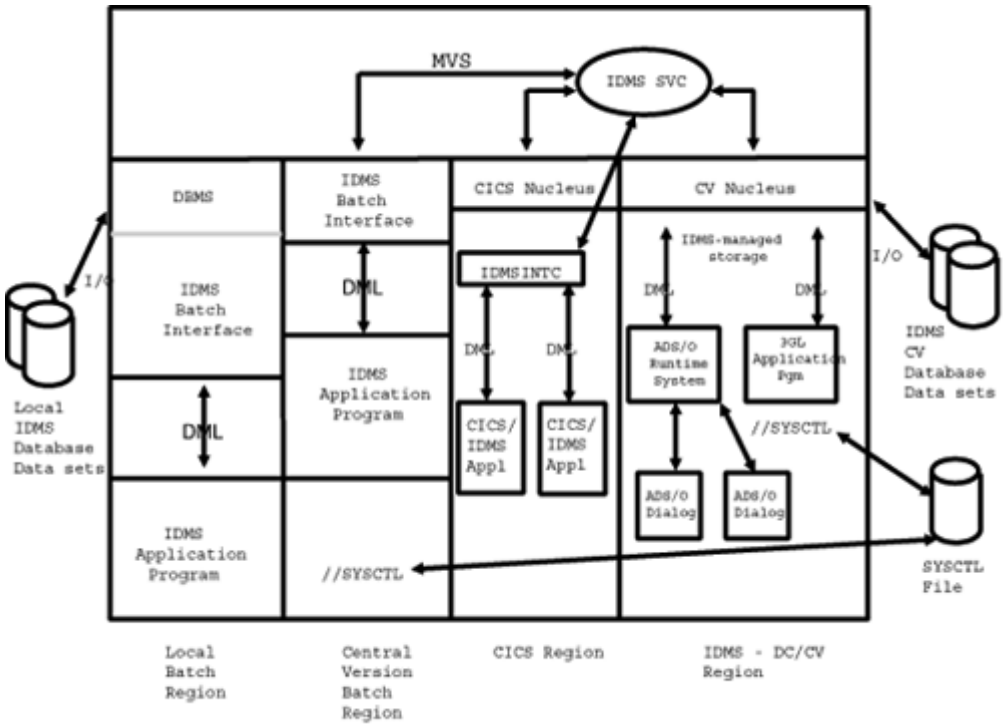

CA-IDMS-DC system services allow the application to:

- GET STORAGE or FREE STORAGE
- LOAD, DELETE, and TRANSFER CONTROL to or RETURN from other application programs
- WAIT on events, POST events complete
- ENQUEUE and DEQUEUE resources
- schedule time-dependent or threshold events (GET TIME, SET TIME)

- communicate with terminal networks (READ TERMINAL, WRITE TERMINAL, MAP IN, MAP OUT)
- ATTACH new tasks and CHANGE PRIORITY of existing tasks
- queue data for use by the same task or other tasks (GET, PUT, and DELETE QUEUE)
- issue messages and acquire CA-IDMS system information (#WTL, #ACCEPT)

Additionally, CA-IDMS-DC manages the terminal network, controls the task dispatching process, and uses service tasks as necessary to support system facilities such as terminal paging.

CA-IDMS-CV database services allow the application to OBTAIN, MODIFY, STORE, ERASE, and KEEP database records. During request processing, CA-IDMS-CV writes to the CA-IDMS journal files images of database records before and after update. Applications can programmatically choose to COMMIT or ROLLBACK changes made to the database.

These services provide a basis for the development of sophisticated online, real-time applications written in assembler, COBOL, or PL/I. Applications can be either conversational or pseudo-conversational.

Two reports in the STROBE Performance Profile can help you isolate the control sections responsible for a majority of the CPU activity. The Transaction Usage by Control Section report identifies which control sections within each transaction consumed the most CPU time, and the Program Usage by Procedure report details each control section so that you can determine the exact area of code that is executing. You can use that information to isolate the exact statement that caused the activity and then modify the code to make it more efficient. Once the code has been modified, you can rerun the measurement session and compare the resulting Performance Profiles to determine whether the change was effective.

CA-IDMS-DC online applications are often created using ADS/O. It is an application development system and fourth-generation language that facilitates rapid development through using an online-based approach and automating functions such as 3270 panel processing and pseudo-conversational programming.

The STROBE CA-IDMS Feature identifies ADS/O dialogs, online programs, and external run units that cause CPU activity. STROBE reports on ADS/O dialogs in the Dialog CPU Summary and details activity for each dialog version in the CPU Usage by ADS/O Statement reports. Detailed information is broken down by statement, process version, dictionary name, and dictionary node so that the ADS/O application developer can quickly locate the appropriate dialog using the summary and then use the detailed reports to isolate the processes and statements that consumed the most CPU time. On the Attribution of CPU Execution Time report, STROBE identifies the ADS/O dialog's transaction, dialog name, location, process name, and version that is consuming CPU resources.

The Attribution of CPU Execution Time report also identifies online programs and external run units that cause CPU activity. STROBE identifies activity caused by online programs by their transaction, module, control section, hexadecimal offset, DML sequence number, function, and version. STROBE identifies activity caused by external run units by their batch job name (in the transaction column) program name, DML sequence or SQL line number, and function.

To produce meaningful Dialog reports in the STROBE Performance Profile, you need to have generated diagnostic tables when you compiled your dialogs. If these tables do not exist, STROBE can attribute activity only to an ADS/O FDB CME offset, not to an actual statement number, and it cannot index any of the dialogs.

CA-IDMS Central Version Batch Mode

Database applications can be developed using assembler, COBOL, or PL/I for batch execution. These applications typically handle large amounts of data and run concurrently with online processing. They execute in a separate region and have access to the same Central Version database as the online region. When the application is ready

to begin database access, it signs on to the CV that it will call to satisfy DB DML and SQL requests for database records. The CV fulfills the request and returns the data to the application. All CA-IDMS DB DML commands are available with batch to facilitate database I/O, but only a limited subset of DC DML commands are available since a full data communications environment does not exist.

The STROBE CA-IDMS Feature attributes wait time in CA-IDMS Central Version batch jobs to the database DML or SQL statement responsible for causing the wait. In the Attribution of CPU Wait Time report, the feature pinpoints exactly where developers should focus their improvement efforts by showing the modules, control sections, DML or SQL statement offsets, and, if available, the DML sequence numbers or SQL line numbers, of the statements that caused CPU wait time.

CA-IDMS Local Mode Batch

Database requests are issued from the batch region, a complete single-user database environment allocating and controlling all database files, DBMS programs, and database buffers. All DB DML and SQL requests for database records are satisfied completely within the batch region. There is no dependence on the Central Version online region.

The STROBE CA-IDMS Feature attributes wait time in CA-IDMS local mode batch jobs to the database DML or SQL statement responsible for causing the wait. In the Attribution of CPU Wait Time report, the feature pinpoints exactly where developers should focus their improvement efforts by showing the modules, control sections, DML or SQL statement offsets, and, if available, the DML sequence or SQL line numbers, of the statements that caused CPU wait time.

CA-IDMS-DC/CV-UCF Mode

The CA-IDMS-DC/CV-UCF mode permits the use of CICS (and other teleprocessing monitors) to submit transactions to an active CA-IDMS-DC/CV online region. The transactions are executed by the CV region and the results returned to the CICS region. The terminal user issuing the original transaction on CICS is not aware of where execution occurs. UCF allows an installation that has chosen CICS as its standard teleprocessing monitor, for example, to be able to execute CA-IDMS-based applications, including ADS/O, through CICS.

Additionally, the application developer can issue CA-IDMS DB DML and SQL requests for database records contained in a CA-IDMS Central Version database from within CICS command- or macro-level assembler, COBOL, and PL/I programs.

The STROBE CA-IDMS Feature can be used to measure the CA-IDMS-DC/CV regions where the actual application activity is taking place. If you want to analyze another teleprocessing monitor, such as CICS, that is shipping transactions and DML or SQL requests to a CA-IDMS-DC/CV region, examine the Transaction Summary and the Transaction Usage by Control Section reports. In these reports, you should check

- the transaction names, if you are analyzing a shipped transaction
- the CICS region job names, if you are analyzing a DML request shipped from a CICS application program

If a transaction is shipped to CA-IDMS via UCF for execution, all the transaction's activity occurs within the CA-IDMS-DC/CV region. If a DML or SQL request is shipped by a CICS application program, the CICS region is acting as an external requestor of database services. For more information, see the section titled “Transaction Reports” on page 3-11.

Identification of Activity in Load Modules

STROBE identifies activity within CA-IDMS system modules — load modules that have names beginning with IDMS or RHDC — to the pseudo-section .IDMS within the pseudo-module .SYSTEM. The true name appears as a procedure within .IDMS. The STROBE CA-IDMS Feature provides function descriptors for the CA-IDMS modules.

To obtain detailed reporting on a CA-IDMS module, specify its name with the DETAIL parameter of the STROBE Command Language. For a detailed explanation of the reporting options, see Chapter 3 of the *STROBE MVS User's Guide* or *STROBE MVS User's Guide with Advanced Session Management*.

Identification of Activity in Taskcodes

CA-IDMS system tasks are identified with a pseudo-taskcode that indicates the function of the task. The following table lists the STROBE pseudo-taskcodes and the tasks they represent.

CA-IDMS Pseudo-Taskcodes (Transactions)

Pseudo-name	Description
LDRVR	CA-IDMS DC Line Driver Task
PDRVR	CA-IDMS DC Printer Driver Task
DBRC	CA-IDMS DC Database Recovery Control Task
ERUS	CA-IDMS DC External Run-unit Service Task
FACTOT	CA-IDMS DC Factotum Task
HELOT	CA-IDMS DC Helot Task
MSTR	CA-IDMS DC Master Task

STROBE-Generated Module Names

The STROBE-generated module name IDMSXXX represents activity within CA-IDMS-managed storage areas that cannot be attributed to a specific load module.

The STROBE-generated module name IDMSSVC represents activity within a CA-IDMS-managed control block that issues SVC calls to the CA-IDMS SVC.

Other Pseudo-Names

In some cases, neither the STROBE sampling routines nor the STROBE CA-IDMS Feature can identify the executing CA-IDMS or MVS module. In these cases, STROBE attributes execution to a pseudo-section according to the location of the executable code within the address space. (For more information, see Appendix A of *STROBE MVS Concepts and Facilities*.)

Benefits of the STROBE CA-IDMS Feature

The STROBE CA-IDMS Feature will help you throughout the application life cycle:

- in design and build, to assess alternative programming techniques by measuring the performance of various test cases
- in test and quality assurance, to confirm design assumptions about new or modified application resource requirements
- in production to ensure the performance of ADS/O and CA-IDMS batch programs and usage of CA-IDMS services
- in maintenance, to assess the impact of changes in business requirements and data volume on the performance of ADS/O and CA-IDMS batch applications and the CA-IDMS database management.

Chapter 2.
Using the STROBE CA-IDMS Feature

This chapter covers some general information you need to start and configure the CA-IDMS Feature. It also explains how to index ADS/O dialogs for the STROBE Performance Profile. Sections describing how to create the indexes using either ISPF panels or the STROBE command language are included.

Invoking the STROBE CA-IDMS Feature

Depending on whether you are measuring an online region or a batch job, STROBE uses different ways to invoke the CA-IDMS Feature. This section briefly explains the differences. You can find detailed instructions for submitting measurement requests in the STROBE/ISPF Online Tutorials and in Chapters 2 and 3 of the *STROBE MVS User's Guide* or the *STROBE MVS User's Guide with Advanced Session Management*.

CA-IDMS-DC/CV Online Regions

If you want measurement data (including ADS/O attribution data) from a CA-IDMS-DC/CV online region, STROBE automatically invokes the STROBE CA-IDMS Feature when

- you measure a target program whose name (as given in the PGM= operand of the EXEC statement that invokes the job step) begins with "IDMS"
- your STROBE system programmer specified an appropriate alias to STROBE during installation

Otherwise, you must specify that STROBE should start the STROBE CA-IDMS Feature when you add a measurement request by doing the following:

- with STROBE/ISPF, enter "Y" in the DATA COLLECTORS field of the STROBE - ADD ACTIVE or STROBE - ADD QUEUED panel, then specify "Y" in the IDMS field on the DATA COLLECTORS panel (Figure 2-1 on page 2-2)
- with the STROBE command language, include the IDMS keyword parameter when you submit an ADD operation

If you have no interest in CA-IDMS measurement data, you can stop STROBE from automatically invoking the data collector by doing the following:

- with STROBE/ISPF, enter "Y" in the DATA COLLECTORS field of the STROBE - ADD ACTIVE or STROBE - ADD QUEUED panel, then specify "N" in the IDMS field on the DATA COLLECTOR panel
- with the STROBE command language, include the NOIDMS keyword parameter when you submit an ADD operation

Note: You *cannot* collect attribution data unless the STROBE CA-IDMS Feature is active. Specifying NOIDMS disables all CA-IDMS data collection.

If you have no interest in ADS/O attribution data that is available in the Dialog CPU Summary and CPU Usage by ADS/O Statement reports, you can suppress ADS/O attribution data collection by doing the following:

- with STROBE/ISPF, enter “Y” in the OTHER PARAMETERS field of the STROBE - ADD ACTIVE or STROBE - ADD QUEUED panel, then specify on the OTHER PARAMETERS panel

  ```
  IDMS=NOADSO
  ```

- with the STROBE command language, submit an ADD operation and include the operan yod

  ```
  IDMS=NOADSO
  ```

Figure 2-1. STROBE - DATA COLLECTORS Panel

```
--------------------------- STROBE - DATA COLLECTORS --------------------------
COMMAND ===>

OVERRIDE DATA COLLECTOR DEFAULTS FOR JOBNAME:  WPAFXC

DATA COLLECTORS: (Y or N; Y adds to and N removes from your system defaults)
  ADABAS          ===>     ADA3GL          ===>     C                ===>
  CICS            ===>     COBOL           ===>     CSP              ===>
  DB2             ===>     IDMS            ===> Y   IDMS BATCH DML   ===>
  IEF             ===>     IMS             ===>     NATURAL          ===>
  PL/I            ===>     SVC             ===>


CICS Options: (Y or N; default is Y)
 Produce Performance Supplement ===>   Collect terminal activity ===>


CAPTURE Options: (Y or N; default is Y)
          DB2    ===>    IMS    ===>


OTHER DATA COLLECTORS:
  PROGRAM NAME   ===>           ===>           ===>           ===>
  PROGRAM NAME   ===>           ===>           ===>           ===>
```

CA-IDMS Central Version and Local Mode Batch Regions

STROBE does not automatically invoke DML attribution when you measure a batch job that issues database DML requests. To obtain DML attribution, you must specify the feature. To specify the DML feature do the following:

- With STROBE/ISPF, enter “Y” in the DATA COLLECTOR field of the STROBE - ADD ACTIVE or STROBE - ADD QUEUED panel, then enter “Y” in the IDMS BATCH DML field shown previously in Figure 2-1.
- With STROBE command language, specify DML.

Note: If you specify DML attribution, then the STROBE CA-IDMS Feature is automatically invoked. Do not both specify DML attribution and enter “Y” next to IDMS on the STROBE - DATA COLLECTORS panel. When using STROBE command language, do not specify IDMS if you have specified DML.

To show wait related to DML sequence numbers, you must compile the program with the DEBUG DML option selected.

Indexing ADS/O Dialogs

When the STROBE CA-IDMS Feature measures activity within dialogs that were generated with diagnostic tables, it can obtain the statement numbers of the ADS/O commands that are executing. These dialogs can then be indexed, allowing STROBE to associate an ADS/O statement number with its corresponding source statement text. When a dialog

has been indexed, STROBE includes both the statement number and all or part of the source statement text in the CPU Usage by ADS/O Statement report. Otherwise, STROBE displays only the statement numbers (if diagnostic tables exist) or the ADS/O FDB CME offsets (if they do not).

You can index your dialogs by executing a separate batch job any time before creating the STROBE Performance Profile. When you request indexing, specify the name of the listing data sets created by the CA-IDMS ADSORPTS utility program, as described below.

Unless you generate diagnostic tables when you compiled your dialogs, STROBE can attribute activity only to an ADS/O FDB CME offset, not to an actual statement number, and it cannot index any of the dialogs. If your Performance Profile displays only the CME offsets, regenerate the dialogs with diagnostic tables and then rerun the measurement session so you can index the dialog.

Producing ADS/O Dialog Listing Data Sets

Before you can execute the STROBE ADS/O Indexer STRXADS, you first have to generate the listing data set STROBE uses to match the dialog measurement data it has gathered with the corresponding ADS/O statements.

To generate a listing data set containing the source code for each ADS/O dialog, run the CA-IDMS utility program ADSORPTS with the run-time parameter REPORTS=ALL. (See Chapter 4 of the *STROBE MVS User's Guide* or *STROBE MVS User's Guide with Advanced Session Management* for a list of the DCB parameters you need to set to create the listing data set.) To specify multiple dialogs, use the ADSORPTS DIALOGS parameter, which accepts any of the following values:

- ALL (generate reports for all dialogs).
- Dialog-name (generate a report for a single dialog). To specify multiple dialog names, enter them in a comma-delimited list enclosed by parentheses.
- Dialog-mask-a (generate a report for any dialog that matches the mask criteria).
- Low-dialog-name-high-dialog-name (generate a report for any dialogs within a specified range). The hyphen is required and cannot have surrounding blanks.

The following job control statements will produce reports for all versions of the dialogs named DIAG1, DIAG2, and DIAG3. The ADS/O listing will be placed in the data set named SA.STDM.ADSORPTS.DIAG1.

```
//ADSORPTS EXEC PGM=ADSORPTS,REGION=2048K
//STEPLIB  DD DSN=SA.STDM.IDMS.LOADLIB,DISP=SHR
//CDMSLIB  DD DSN=SA.STDM.IDMS.LOADLIB,DISP=SHR
//TDICT    DD DSN=SA.STDM.IDMS.TDICT,DISP=SHR
//TDLOD    DD DSN=SA.STDM.IDMS.TDLOD,DISP=SHR
//SYSPCH   DD DUMMY
//SYSJRNL  DD DUMMY
//SYSLST   DD DSN=SA.STDM.ADSORPTS.DIAG1,DISP=SHR
//SYSIPT   DD *
 DBNAME=DICTTD
 LIST=FULL
 DIALOGS=(DIAG1,DIAG2,DIAG3),VERSIONS=ALL,REPORTS=ALL
```

After you have created the listing data set, supply its name to the STROBE Indexer STRXADS, either through STROBE/ISPF or with the STROX procedure, as described in the next section.

An ADS/O Indexer map data set is identical in format to those produced by the other STROBE Indexers. When you produce a Performance Profile, you can concatenate ADS/O Indexer map data sets with one another and with map data sets produced by the other STROBE Indexers.

Using STROBE/ISPF to Index ADS/O Dialogs

To index your ADS/O dialogs with STROBE/ISPF, select Option 5, INDEX, from the STROBE OPTIONS menu. Then enter the following values on the INDEX TO CREATE A MAP DATA SET panel (Figure 2-2):

- select background or foreground processing by entering the appropriate code in the OPTION field
- enter “Y” in the ADS/O field
- specify the name of the map data set you want to create in the OUTPUT: MAP DATA SET field
- specify the name of the listing file produced by the ADSORPTS utility program in the first INPUT: COMPILER SYSPRINT DATA SETS field

For example, if you want to submit a background processing request to create a map data set called SA.STDM.DIAG1.MAP from the ADSORPTS listing in SA.STDM.ADSORPTS.DIAG1, you would enter the information on the INDEX TO CREATE A MAP DATA SET panel as shown in Figure 2-2.

Figure 2-2. STROBE - INDEX TO CREATE A MAP DATA SET Panel

```
------------------- STROBE - INDEX TO CREATE A MAP DATA SET -------------------
OPTION ===> B
            B  - Background processing   F  - Foreground processing

PROGRAM LANGUAGE: (Specify Y -- select one language)
 ADABAS/NATURAL  ===>       ADS/O           ===> Y     ASSEMBLER       ===>
 C (IBM OR SAS)  ===>       CA OPTIMIZER    ===>       COBOL           ===>
 CSP             ===>       DB2 DBRM        ===>       FORTRAN G       ===>
 FORTRAN VS OR H ===>       IEF             ===>       PL/I            ===>

OUTPUT: MAP DATA SET
       ===> 'sa.stdm.diag1.map'
  UNIT ===> WPAANY     VOLUME ===>

INPUT: COMPILER SYSPRINT DATA SETS
       ===> 'sa.stdm.adsorpts.diag1'
       ===>
       ===>
       ===>
       ===>
       ===>
```

Using the STROX Procedure to Index ADS/O Dialogs

You can execute the STROBE procedure STROX to index the program and create a map data set, supplying these parameters:

- MAPPGM — STRXADS (the name of the Indexer program)
- LISTDSN — the output file name produced by the ADSORPTS utility program
- MAPDSN — the name of the output map data set

```
//INDEX    EXEC  STROX,
//               MAPPGM=STRXADS,
//               LISTDSN='SA.STDM.ADSORPTS.DIAG1',
//               MAPDSN='SA.STDM.DIAG1.MAP'
```

To concatenate multiple reports from the ADSORPTS utility and produce only one map data set, override the LISTING DD statement in the STROX procedure with a concatenation of the ADSORPTS listings data set names.

```
//INDEX    EXEC STROX,
//               MAPPGM=STRXADS,
//               MAPDSN='SA.STDM.DIAG1.MAP'
//LISTING   DD   DISP=SHR,DSN=SA.STDM.ADSORPTS.DIAG1
//               DISP=SHR,DSN=SA.STDM.ADSORPTS.DIAG2
```

Using Your Own Job Control Statements

Alternatively, you can code the job control statements to execute the CA-IDMS Indexer as shown below:

```
//INDEX    EXEC  PGM=STRXADS
//STEPLIB   DD   DISP=SHR,DSN=STROBLIB
//SYSPRINT  DD   SYSOUT=*
//STRMAP    DD   DISP=SHR,DSN=SA.STDM.DIAG1.MAP
//LISTING   DD   DISP=SHR,DSN=SA.STDM.ADSORPTS.DIAG1
```

Tailoring STROBE Reports

This section explains how to tailor reports in a Performance Profile so only specific measurement data you are seeking is provided. For a complete description of creating STROBE Performance Profile reports, see Chapter 4 of the *STROBE MVS User's Guide*.

Suppressing Attribution Reports in the Performance Profile

To limit or suppress the Attribution of CPU Execution Time and Attribution of CPU Wait Time reports in the Performance Profile, you can specify any of the parameters below.

Suppressing Attribution Reports with Low Activity

You can suppress Attribution reports for all system service modules in which the total CPU or wait time percentage is less than a specified baseline by doing the following:

- With STROBE/ISPF, enter "Y" in the Tailor Reports field of the STROBE - PRODUCE A PERFORMANCE PROFILE panel, then, on the STROBE - TAILOR REPORTS panel (Figure 2-3 on page 2-6), specify a baseline percentage between 0 and 99.9 in the Compress below % column of the ATTRIBUTION Reports field.
- When you submit a batch job using the STROE or STROXE procedure, include ATTR=*nn.n*, where *nn.n* is the baseline percentage (default 2).

Figure 2-3. STROBE - TAILOR REPORTS Panel

```
-------------------------- STROBE - TAILOR REPORTS --------------------------
COMMAND ===>

WAIT TIME BY MODULE -- Show location of wait  ===>            (Specify Y for a
                                                              selection except
                                     Compress                 where % required)
-------------- Report -------------- below %   or  Suppress
PROGRAM USAGE BY PROCEDURE            ===>          ===>
DASD USAGE BY CYLINDER                ===>          ===>
TRANSACTION USAGE BY CONTROL SECTION  ===>          ===>
ATTRIBUTION Reports                   ===> 2.5      ===>

  Suppress reports for C         ===>   CICS      ===>   COBOL     ===>
                        CSP      ===>   DB2       ===>   DL/I      ===>
                        IDMS     ===>   IEF       ===>   PL/I      ===>
                        SVC      ===>

PROGRAM SECTION USAGE SUMMARY             Display inactive  ===>
TIME and RESOURCE DEMAND DISTRIBUTION     Combine tasks     ===>
                                          Display all tasks ===>
                                          Display all DDs   ===>


USE DATE AND TIME FORMAT FROM PARMLIB ===>
```

Suppressing All Attribution Reports

To suppress all the Attribution reports:

- with STROBE/ISPF, enter “Y” in the Tailor Reports field of the STROBE - PRODUCE A PERFORMANCE PROFILE panel, then enter “Y” in the Suppress column of the ATTRIBUTION Reports field on the STROBE - TAILOR REPORTS panel
- with the STROE or STROXE procedure, include the NOATTR parameter

Suppressing CA-IDMS-Specific Attribution Reports

To suppress the Attribution of CPU Execution Time and Wait Time reports for CA-IDMS ADS/O dialogs and batch database application programs only, you can:

- with STROBE/ISPF, enter “Y” in the Tailor Reports field of the STROBE - PRODUCE A PERFORMANCE PROFILE panel, then enter “Y” in the Suppress reports for IDMS field on the TAILOR REPORTS panel
- with the STROE or STROXE procedure, include the NOATTR=IDMS parameter

Chapter 3.
The Performance Profile for Online Regions

This chapter describes the STROBE Performance Profile produced when you measure a CA-IDMS-DC/CV online region and presents a guide to interpreting the reports by showing selections from a Performance Profile. You can find a complete description of the Performance Profile in Chapter 4 of the STROBE *STROBE MVS Concepts and Facilities*.

Measurement Session Data Report

The first report in the Performance Profile, the Measurement Session Data report (Figure 3-1), describes the environment during a measurement session and provides some measurement statistics. When STROBE measures a job step that uses CA-IDMS, the report shows the IDMS release number in the SUBSYSTEM field.

For CA-IDMS Releases 12.0, 12.01, and 14.0 the format is “IDMS 12.0“, “IDMS 12.01“, or “IDMS 14.0“. For CA-IDMS Release 10.2, the format is “IDMS Snnnnn”, where nnnnn is the release number. The release number is followed by a three-character code indicating whether the address space is running Central Version batch (code “CVB”) or local mode (code “LOC”). No characters appear if the measurement session is for an online region. Figure 3-1 is an example of a Measurement Session Data report from a CA-IDMS 12.0 measurement.

Figure 3-1. Measurement Session Data Report

```
                                        ** MEASUREMENT SESSION DATA **

------- JOB ENVIRONMENT --------        ------- MEASUREMENT PARAMETERS ------      ----- MEASUREMENT STATISTICS ----

PROGRAM MEASURED   -    IDMSPROD        ESTIMATED SESSION TIME  -       5 MIN       CPS TIME PERCENT  -         35.56
JOB NAME           -    IDMSPROD        TARGET SAMPLE SIZE      -      10,000       WAIT TIME PERCENT -         64.44
JOB NUMBER         -    STC08056        REQUEST NUMBER   (A)    -          46       RUN MARGIN OF ERROR PCT -     .62
STEP NAME   -           IDMSPROD        FINAL SESSION ACTION    -        QUIT       CPU MARGIN OF ERROR PCT -     .33
                                                                                    TOTAL SAMPLES TAKEN -      10,000
DATE OF SESSION    -    06/23/96        OPTIONS       -   IDMS   ADS/O             TOTAL SAMPLES PROCESSED -  10,000
TIME OF SESSION    -    14.05.46        MODULE MAPPING BASELINE  -          0       INITIAL SAMPLING RATE - 83.33/SEC
CONDITION CODE     -      C-0000                                                    FINAL SAMPLING RATE   - 83.33/SEC

SYSTEM -             ESA SP4.2.2         --------- REPORT PARAMETERS --------      SESSION TIME -    9 MIN 36.74 SEC
DFSMS        -             1.1.0                                                   CPU TIME          2 MIN 21.46 SEC
SUBSYSTEM   -        IDMS   12.0        REPORT RESOLUTION    -       64 BYTES      WAIT TIME    -    5 MIN 46.18 SEC
CPU MODEL          -   3090-600S        SORTSIZE         -           999,999       STRETCH TIME -    1 MIN 29.10 SEC
SMF/SYSTEM ID      -        IAS6        LINES/PAGE           -             60
                                                                                    SRB TIME          0 MIN  3.11 SEC
REGION SIZE BELOW 16M -   7,564K                                                    SERVICE UNITS                 37
REGION SIZE ABOVE  -     98,304K        DASD=2.0%   DASDGAP=5                      PAGES IN-    555     OUT-       3
                                        ATTR= 0.0%                                  PAGING RATE  -           0.97/SEC
PTF LVL- 2.5.0.FS000000/FS000000                                                    EXCPS -      1,754       3.05/SEC
STROBE TAPE NUMBER -  000-S85999

SAMPLE DATA SET  - STROBE.SAMPLE.IDMSPROD.S015D001
```

Choosing Between the Execution and Wait Reports

To determine whether to examine the CPU execution or wait reports in the Performance Profile, look under the “MEASUREMENT STATISTICS“ column at the values reported for the “CPS TIME PERCENT“ (the percentage of time that one or more CPUs were active) and WAIT TIME PERCENT fields. If the value for WAIT TIME PERCENT is very high, then you should examine the Wait Time by Module and Attribution of CPU Wait Time reports

to determine which programs and control sections are waiting. For more information on wait time, see the section called "Wait Time by Module Report" on page 3-15. If the value for CPS TIME PERCENT is high, see "Reports That Show CPU Execution" below.

In Figure 3-1 on page 3-1, the CPS TIME PERCENT value of 35.56% suggests that performance might be improved by reducing CPU usage. The WAIT TIME PERCENT value of 64.44 indicates that trying to curtail wait time might also improve performance. First you should examine the CPU execution reports and then analyze the wait time reports.

Reports That Show CPU Execution

The Performance Profile shows CPU execution in ADS/O and batch regions in four types of reports. Each report type shows a different view of CPU execution activity during a measurement session.

Program Module Reports

The Program Section Usage Summary and Program Usage by Procedure reports show CPU usage in all programs, including application programs, CA-IDMS service modules, and other system service routines. You can use the Program Section Usage Summary to determine which CA-IDMS service modules or user programs were responsible for the majority of the CPU usage. The Program Usage by Procedure reports detail CPU usage by user programs and system service modules.

The Most Intensively Executed Procedures report highlights the information in the module reports. It lists the ten heaviest users of CPU time and lets you quickly determine whether service modules or user programs are responsible for a majority of the CPU usage.

ADS/O Dialog Reports

The Dialog CPU Summary and the CPU Usage by ADS/O Statement reports show which dialogs and, within a dialog, which processes and ADS/O statements are responsible for CPU activity. Turn to these reports when most of the CPU activity occurs in dialogs.

Attribution Reports

The Attribution of CPU Execution Time report identifies which ADS/O dialogs, online programs, and external run units are callers of CA-IDMS system service routines. ADS/O dialog callers are identified by transaction (taskcode) name, dialog name, process name and version number, and by either ADS/O statement number or FDB CME offset. When a specific service routine has a high CPU usage, its Attribution of CPU Execution Time report pinpoints the ADS/O statements, programs or external run units that invoked it.

The Attribution of CPU Execution Time identifies programs that consume CPU by transaction, module name, control section name, hexadecimal offset, DML sequence number, function, and version. It identifies CPU activity in external run units by batch job name (in the transaction column), program name, DML sequence number, and function.

Note: To include a DML sequence number, the program must be compiled with the DEBUG DML compiler option.

Transaction Reports

The Transaction Summary and Transaction Usage by Control Section reports identify, for online address spaces, which transactions (taskcodes) are responsible for CPU activity. They show you which transactions were active during the measurement session, and

which transactions consumed the most CPU time. The Transaction Summary report also gives a count of the total number of times the transaction executed and the average service time, which is the average wall clock time between transaction start and transaction end, for each transaction.

Program Module Reports

This section describes the program module reports that provide data STROBE collects about CPU usage during a measurement session.

Program Section Usage Summary Report

The Program Section Usage Summary report (Figure 3-2) shows the distribution of CPU time used by each active control section of each active module in the CA-IDMS address space. CA-IDMS and other system modules are compressed and shown as pseudo-sections under the pseudo-module .SYSTEM.

This report can help to focus your investigation by showing you whether CA-IDMS system modules or user programs are responsible for the most CPU activity. If CA-IDMS system modules and the ADS/O run-time system are the heaviest CPU users, refer to the ADS/O Dialog reports to identify performance improvement opportunities. If user-written DB/DC applications show high activity, refer to the Program Usage by Procedure reports.

In this example, because CA-IDMS and ADS/O modules account for most of the execution time, you should refer to the Dialog reports to determine what caused this activity.

Figure 3-2. Program Section Usage Summary Report

```
                                        ** PROGRAM SECTION USAGE SUMMARY **

 MODULE  SECTION  16M SECTION  FUNCTION            CPU TIME PERCENT       CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
   NAME     NAME  <,>    SIZE                       SOLO     TOTAL        .00     16.50     33.00     49.50     66.00

.SYSTEM  .IDMS                IDMS SYSTEM SERVICES      62.14     68.82        .*************************+++
.SYSTEM  .IOCS                DATA MANAGEMENT SERVICES    .07       .14        .
.SYSTEM  .NUCLEUS             MVS NUCLEUS                3.81      4.09        .**
.SYSTEM  .SVC                 SUPERVISOR CONTROL         6.89      7.39        .****+
.SYSTEM  .VSAM                VIRTUAL STORAGE ACC METH    .22       .22        .
                                                         -----     -----
.SYSTEM  TOTALS               SYSTEM SERVICES           74.00     81.53

ACCTINT           >     16080  INTEREST CALCULATION RTN    .29       .29        .
ACCTM100          >     24600  ACCOUNT MAINTENANCE         .36       .36        .
ACCTM200          >     18200  ACCOUNT MAINTENANCE         .36       .36        .
ADSOGEN2          >    151656  ADS/O GENERATOR ROUTINE     .07       .07        .

ADSOMAIN ADSOCCTL       11804  RUNTIME CNTL CMD INTERP     .07       .14        .
ADSOMAIN ADSOCDAB        7730  RUNTIME DB CMD INTERP      1.80      1.80        .*
ADSOMAIN ADSOCDRV       13105  RUNTIME COMMAND DRIVER     3.01      3.52        .**+
ADSOMAIN ADSOCEXC        3872  RUNTIME EXEC CODE MGR       .36       .43        .
ADSOMAIN ADSOCMAP        3148  RUNTIME OLM CMD INTERP      .07       .07        .
ADSOMAIN ADSOCPRO        7808  RUNTIME PROCEDURAL CMDS    2.51      2.73        .*
ADSOMAIN ADSOCSAI          20  RUNTIME CSA INTERFACE       .07       .07        .
ADSOMAIN ADSOFPRO        4000  RUNTIME CMD FAST PROCESS    .65       .72        .
ADSOMAIN ADSOPTCH        4096  RUNTIME PATCH MODULE        .14       .14        .
ADSOMAIN ADSORUN2       15730  RUNTIME PSEUDO-CONVERSER   3.30      3.30        .**
ADSOMAIN ADSORUN3        8590  RUNTIME COMMON ROUTINES    3.24      3.31        .**
                                                         -----     -----
ADSOMAIN TOTALS    >    82137  RUNTIME DIALOG EXECUTION  15.22     16.23

ADSORUN1          >     7696   RUNTIME ENVIRONMNT SETUP    .14       .21        .
PMVECTRS          >    15872   PERFORMANCE MONITOR        1.44      1.51        .
TSSM102           >    17960   TOP SECRET                  .29       .29        .
                                                         -----     -----
PROGRAM IDMSPROD TOTALS                                  91.31    100.00
```

Program Usage by Procedure

The Program Usage by Procedure reports present a detailed accounting of the CPU time spent by each control section within each program that was active during a measurement session. There are reports for both application programs and system modules, and the two types of reports look slightly different. For system modules, STROBE displays the module and control section name and a function descriptor, as shown in Figure 3-3 on page 3-5.

For application programs, STROBE includes the module name, control section name, and starting location. If the control section was indexed, STROBE shows the line number and procedure name as well. Both types of reports show the solo and total CPU time used by each control section within the module. (For a complete description of this report, see Chapter 3 in *STROBE MVS Concepts and Facilities*.)

Activity in System Modules

The first type of Program Usage by Procedure report details activity in system modules. The example in Figure 3-3 on page 3-5 shows the CA-IDMS modules that support the ADS/O programs and other 3GL applications in this environment.

Of particular interest in this example are module IDMSDBMS (the CA-IDMS database manager), which accounts for 16.16% of CPU activity; module RHDCEVAL (the run-time expression evaluator), which accounts for 6.90%; and module IDMSKEEP (the database key locking routine), which accounts for 6.32%. To determine which transactions and dialogs caused the activity, examine the Attribution reports for these modules.

Activity in Application Programs

The second type of Program Usage by Procedure report shows CPU activity in application programs. This report does not show the CPU activity caused by ADS/O dialogs. Since the ADS/O modules shown in the Program Section Usage Summary report showed significant CPU activity, you should examine the Dialog CPU Summary report (Figure 3-4 on page 3-6) instead of the Program Usage by Procedure report for application programs.

Figure 3-3. Program Usage by Procedure Report fo CA-IDMS Modules

```
                                        ** PROGRAM USAGE BY PROCEDURE **


         .SYSTEM   SYSTEM SERVICES                        .IDMS     IDMS SYSTEM SERVICES

MODULE   SECTION        FUNCTION            INTERVAL  CPU TIME PERCENT       CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
  NAME      NAME                              LENGTH   SOLO     TOTAL        .00      4.50      9.00     13.50     18.00

IDMSDBIO IDMSDBIO  IDMS DB I/O DRIVER          47608   3.95      4.45        .********+
IDMSDBMS           IDMS DB MANAGER            105336  14.15     16.16        .**********************************++++
IDMSDCOM           IDMS RECORD DECOMPRESS        976    .07       .07        .
IDMSINTD           IDMS/DBMS INTERFACE            64   1.08      1.29        .**
IDMSKEEP           IDMS DB KEY LOCKING RTN        64   5.39      6.32        .***********+++
IDMSPROD BLDLLOAD  OS BLDL/LOAD ROUTINES        3412    .07       .07        .
IDMSPROD OS00MISC  MISC OPSYS FUNCTIONS          318    .29       .29        .
IDMSPROD TIMING    CPU TIME ROUTINES            2606    .86      1.08        .*+
IDMSPROD WAITPOST  OS WAIT/POST ROUTINES        3436    .14       .29        .
IDMSRDYX           READY/READY EXTENT RTN         64    .29       .29        .
IDMSSERV           EXTERNAL REQ UNIT DBRC         64    .07       .07        .
IDMSTASK           IDMS EXTERNAL RUN-UNIT         64    .14       .14        .
RHDCAEDT           OLM RUNTIME AUTO EDIT          64   1.08      1.22        .**
RHDCCSA            IDMS COMMON SYSTEM AREA        64   4.06      4.20        .********+
RHDCCURS           CURRENCY SAVE/RESTORE          64    .93      1.01        .**
RHDCDBCO           DB COORDINATION ROUTINE        64    .00       .07        .
RHDCD04W           IDMS CONSOLE LINE DRIVER    70903   1.36      1.36        .***
RHDCEDQ            RESOURCE ENQUEUE ROUTINE       64    .36       .36        .
RHDCEVAL           RUNTIME EXP EVALUATION         64   6.75      6.90        .***************
RHDCEV01           EVAL STRING PROCESSOR        5176    .14       .14        .
RHDCLIMT           RESOURCE LIMIT CHECK RTN       64    .14       .14        .
RHDCLOCK           LOCK/ILE MANAGER               64    .07       .07        .
RHDCLODR           PROGRAM LOAD MANAGER           64   1.72      1.87        .***+
RHDCMAPR           OLM RUNTIME SERVICES           64   1.65      1.65        .***
RHDCMISC           IDMS DC MISC SERVICES          64   1.65      1.72        .***
RHDCMSTR           MASTER-TASK CNTL PRGM          64    .29       .29        .
RHDCPCTL           PROGRAM CONTROL SERVICES       64    .07       .07        .
RHDCQUEM           QUEUE MANAGER                  64    .07       .07        .
RHDCRMGR           RESOURCE CHAIN MGR             64   1.44      1.51        .***
RHDCRUAL           RUN-UNIT MANAGER               64    .07       .07        .
RHDCSCRM           SCRATCH MANAGER                64    .22       .22        .
RHDCSCTY           SECURITY/SIGNON PROCESS        64    .14       .22        .
RHDCSTGP           IDMS DC STORAGE MANAGER        64   4.31      5.10        .*********++
RHDCTIMP           IDMS DC TIME MGMT PGM          64    .07       .07        .
RHDCTIRH           TERMINAL BASIC MODE I/O        64    .07       .07        .
RHDCTSKC           TASK CREATION ROUTINE          64    .07       .07        .
RHDCTSKI           TASK INITIALIZATION RTN        64    .07       .14        .
RHDCWAIT           MULTI-WAIT DISPATCHER          64   4.56      5.07        .**********+
                                                       -----     -----
         .IDMS     TOTALS                              62.14     68.82
```

Dialog Reports

These reports show the CPU activity in CA-IDMS, and in other system service routines that can be directly attributed to an ADS/O dialog, as a percentage of all CPU activity in the measurement session.

Dialog CPU Summary Report

The Dialog CPU Summary report (Figure 3-4 on page 3-6) shows the distribution of CPU activity among the dialogs that invoked system service routines during the measurement session.

STROBE identifies dialogs by their name, version number, and the dictionary name and dictionary node from which they were called. The execution count gives the number of times that the dialog actually executed. The CPU time given is based on STROBE sampling and is not directly related to the execution count.

When two or more active dialogs have the same name and version number and were called from the same dictionary name and dictionary node but were created at different dates and times, STROBE distinguishes them by appending a hyphen and an instance number to the name of the second through hundredth instance. When there is more than one instance of a dialog executing, the total execution count will be given at the first instance of the dialog. The dialog's creation date and time does not appear on the

Dialog CPU Summary report but will appear on the subsequent CPU Usage by ADS/O Statement report.

In this example, dialog PZTD2550 was the heaviest user of CPU time, accounting for 10.85% of the total CPU activity.

Figure 3-4. Dialog CPU Summary Report

```
                                        ** DIALOG CPU SUMMARY **

DIALOG NAME  VERSION DICTNAME DICTNODE   EXECUTION            CPU TIME PERCENT CPU TIME HISTOGRAM  MARGIN OF ERROR: 2.63%
                                          COUNT               SOLO     TOTAL  .00      3.00      6.00      9.00    12.00

 PZTD0000    0001     EMPDICT  DEFAULT        11               .79    .79      .**
 PZTD0009    0001     EMPDICT  DEFAULT         5               .07    .07      .
 PZTD0009-1  0001     EMPDICT  DEFAULT                         .14    .14      .
 PZTD0300    0001     EMPDICT  DEFAULT         2               .14    .14      .
 PZTD0350    0001     EMPDICT  DEFAULT        12               .22    .22      .
 PZTD0820    0001     EMPDICT  DEFAULT         7               .43    .43      .*
 PZTD1010    0001     EMPDICT  DEFAULT         2               .07    .07      .
 PZTD1240    0001     EMPDICT  DEFAULT         8               .72    .72      .**
 PZTD1242    0001     EMPDICT  DEFAULT         1               .07    .07      .
 PZTD1300    0001     EMPDICT  DEFAULT         2               .07    .07      .
 PZTD1310    0001     EMPDICT  DEFAULT         7               .43    .43      .*
 PZTD1311    0001     EMPDICT  DEFAULT         1               .07    .07      .
 PZTD1400    0001     EMPDICT  DEFAULT        15               .57    .65      .*+
 PZTD1400-1  0001     EMPDICT  DEFAULT                         .14    .14      .
 PZTD1400-2  0001     EMPDICT  DEFAULT                         .22    .22      .
 PZTD1610    0001     EMPDICT  DEFAULT         1               .07    .07      .
 PZTD1611    0001     EMPDICT  DEFAULT         3               .22    .22      .
 PZTD2500    0001     EMPDICT  DEFAULT         4               .65    .72      .**
 PZTD2550    0001     EMPDICT  DEFAULT        30             10.49  10.85     .*********************************++
 PZTD3492    0001     EMPDICT  DEFAULT         6               .14    .14      .
 PZTD3511    0001     EMPDICT  DEFAULT         2               .00    .14      .
 PZTD3600    0001     EMPDICT  DEFAULT         4               .36    .36      .*
 PZTD3611    0001     EMPDICT  DEFAULT         2               .07    .14      .
 PZTD4710    0001     EMPDICT  DEFAULT         3               .07    .07      .
 PZTD6050    0001     EMPDICT  DEFAULT         2               .07    .07      .
                                                             -----  -----
DIALOG TOTALS                                               16.31  17.03
```

CPU Usage by ADS/O Statement

The CPU Usage by ADS/O Statement report (Figure 3-5 on page 3-7) shows the distribution of CPU activity among ADS/O statements within each dialog process.

In this example, the report shows activity for process PM and process RP-ENTER within dialog PZTD2550. Because the dialog was indexed before the report was generated, STROBE displays both the ADS/O statement numbers and text.

Figure 3-5. CPU Usage by ADS/O Statement Report

```
                                        ** CPU USAGE BY ADS/O STATEMENT **


DIALOG PZTD2550      VER 0001   DICTIONARY NAME-EMPDICT   DICTIONARY NODE- **DEFAULT** CREATED 07/17/97 12:19:48
PROCESS PZTD2550-PM                        VER 0001  TIMES EXECUTED-       30   AVG SERVICE TIME-    0.55

  STMT      STATEMENT                                CPU TIME PERCENT       CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
 NUMBER       TEXT                                   SOLO     TOTAL        .00      1.00      2.00      3.00      4.00

    2700      DISPLAY MSG CODE IS 912665.               .86       .86        .********
    3200      DISPLAY MSG CODE IS 912770.              1.36      1.36        .*************
   11200      DISPLAY MSG CODE IS 921804.               .36       .36        .***
   18000      OBTAIN OWNER WITHIN CLAIMS.               .14       .22        .*+
   18300      IF SET CLMPRTY MEMBER                     .07       .07        .
   19800      MOVE CONCAT (EXTRACT(FIRST-NAME,         2.73      2.87        .***************************+
   20500      OBTAIN OWNER WITHIN CLAIMS.               .43       .43        .****
   23000      MOVE DATEDIF (DATE, ACTUAL-GROUP) TO      .07       .07        .
   24600      MOVE ' ' TO SALARY-STATUS-CODE(SUBR)      .07       .07        .
   29900      DISPLAY MSG CODE IS 912665.               .50       .50        .*****
   35000      FIND CURRENT FEATURE.                     .79       .79        .*******
   38500      OBTAIN NEXT FEATURE WITHIN REP            .22       .29        .**
   53800      ACCEPT DBKEY INTO FIRST-PAGING-DBKEY     1.02      1.02        .**********
   54100      ACCEPT DBKEY INTO LAST-PAGING DBKEY      1.07      1.07        .**********
                                                      -----     -----
DIALOG PZTD2550      PZTD2550-PM                      9.70      9.99


DIALOG PZTD2550      VER 0001   DICTIONARY NAME- EMPDICT  DICTIONARY NODE- **DEFAULT** CREATED 07/17/97 12:19:48
PROCESS PZTD2550-RP-ENTER                     VER 0001   TIMES EXECUTED-        30  AVG SERVICE TIME-    0.75
  STMT      STATEMENT                                CPU TIME PERCENT       CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
 NUMBER       TEXT                                   SOLO     TOTAL        .00       .50      1.00      1.50      2.00

     900      DISPLAY MSG CODE IS 912980.               .07       .07        .*
    5000      FIND CALC CLAIM.                          .14       .22        .**++
   11300      DISPLAY MSG CODE 912502.                  .22       .22        .****
   13500      DISPLAY MSG CODE 912665.                  .14       .14        .**
   15000      FIND FEATURE DB-KEY IS LAST-PAGING-DBKE   .07       .07        .*
   15700      OBTAIN FEATURE DB-KEY IS LAST-PAGING-DB   .07       .07        .*
    2000*1    MOVE   AGR-CURRENT-FUNCTION       TO      .07       .07        .*
                                                      -----     -----
DIALOG PZTD2550      PZTD2550-RP-ENTER                 .79       .86

DIALOG PZTD2550             DIALOG TOTALS            10.49     10.85
```

Dialog Identification

In the header line of each report, STROBE shows the dialog name, its version number, the process module name, its version number, the dictionary name and dictionary node from which the dialog was called, the dialog's creation date and time, the number of times STROBE observed the process module executing, and its average service time. The TIMES EXECUTED and AVERAGE SERVICE TIME are calculated based on the number of times STROBE observes the process module executing. These values are not directly related to the CPU time reported.

Statement Number

If diagnostic tables were requested when the dialog was generated, STROBE displays the statement numbers in the STMT NUMBER field. If the line is part of an INCLUDEd code section, then STROBE appends an asterisk (*) and a number signifying the unique occurrence of the INCLUDE to the statement number. For example, the statement number 2000*1 refers to statement number 2000 in the first INCLUDE code block occurrence in process PZT2550-RP-ENTER. The occurrence number refers to the order of appearance in the dialog process of the INCLUDEd code block in which the statement number is contained.

If diagnostic tables were not available, STROBE displays the ADS/O FDB command (CME) offset in this field.

The statement numbers in the Dialog reports match those in the listing produced by the STROBE CA-IDMS Indexer. This Indexer listing contains the statement number and complete text of every ADS/O statement in a dialog.

Statement Text

If an ADS/O dialog has been indexed, STROBE includes the statement text in the CPU Usage by ADS/O Statement report.

As shown in Figure 3-5 on page 3-7, statements 19800 MOVE CONCAT, 53800 ACCEPT DBKEY, and 54100 ACCEPT DBKEY within process PZTD2550-PM are all relatively heavy users of CPU time. An examination of the code (Figure 3-6) shows that statement 19800 combines a MOVE with CONCAT function and an EXTRACT function. Both of these ADS/O functions are very CPU-intensive, and the combination should be avoided whenever possible. In this case, you can reduce the complexity of the operations that ADS/O must perform by using a series of single MOVE statements in place of the CONCAT function.

Figure 3-6. ADS/O Dialog Code Showing CONCAT

```
00019300      IF LAST-NAME EQ SPACES
00019400          MOVE 'UNKNOWN' TO LAST-NAME.
00019500      IF FIRST-NAME EQ SPACES
00019600          MOVE 'UNKNOWN' TO FIRST-NAME.
00019700
00019800      MOVE CONCAT (EXTRACT(LAST-NAME), ', ',
00019900                   EXTRACT(FIRST-NAME), ' ',
00020000                           MID-INIT, ' ',
00020100                           NAME-SUFXX) TO
00020200                           NAME (MAP-SUB).
00020300
00020400  !***MOVE TEMP-ORIGIN-CODE TO ORIGIN-CODE (ORG-SUB).
00020500      OBTAIN OWNER WITHIN CLAIMS.
```

In the case of statements 53800 and 54100, an examination of their code (Figure 3-7) shows that two ACCEPTs are being used to load the same value into two different fields. Because both statements are relatively CPU-intensive, the ACCEPT statement 54100 should be replaced with MOVE FIRST-PAGING-DBKEY TO LAST-PAGING-DBKEY. Also note that statements 52900 and 53200 exhibit the same characteristics and can benefit from a similar replacement.

Figure 3-7. ADS/O Dialog Code Showing ACCEPTS

```
00052200  !=====================================================================
00052300  DEFINE SUBROUTINE SAVDBKEY.
00052400  !=========================================================================
00052500
00052600  IF CURRENT-RESPONSE = 'BACKWARD'
00052700     DO.
00052800       IF ORG-SUB EQ 14
00052900          ACCEPT DB-KEY INTO LAST-PAGING-DBKEY
00053000                 FROM FEATURE CURRENCY.
00053100
00053200       ACCEPT DB-KEY INTO FIRST-PAGING-DBKEY
00053300                 FROM FEATURE CURRENCY.
00053400     END.
00053500  ELSE
00053600     DO.
00053700       IF ORG-SUB EQ 1
00053800          ACCEPT DB-KEY INTO FIRST-PAGING-DBKEY
00053900                 FROM FEATURE CURRENCY.
00054000
00054100       ACCEPT DB-KEY INTO LAST-PAGING-DBKEY
00054200                 FROM FEATURE CURRENCY.
00054300     END.
00054400
00054500  GOBACK.
00054600
```

Attribution Reports

This section describes the attribution reports that identify ADS/O dialogs, programs, and external run units that call CA-IDMS system service routines.

Attribution of CPU Execution Time

The Attribution of CPU Execution Time report (Figure 3-8 on page 3-10) identifies the sites of invocation of system routines within ADS/O dialogs, programs, and external run units. Use this report to identify which ADS/O statements, programs, and external run units caused the most activity.

Header Lines

The report header identifies the invoked routine, showing

- its pseudo-module, module, and control section name (when available)
- a function descriptor for either the control section or the module

Detail Lines

Detail lines for ADS/O dialogs that call CA-IDMS service routines identify a module invoker and include

- an ADS/O taskcode
- the dialog name (and appended instance number, if any)
- the ADS/O statement location — either the statement number followed by an asterisk and any INCLUDE occurrence of the statement, or the eight-digit CME offset if the dialog was generated without diagnostic tables
- the ADS/O process name
- the process version number
- the solo and total CPU time spent on behalf of the invoker

Detail lines for external run units specify

- in the transaction name column, the batch job name
- the program name
- the DML sequence number
- the function

Detail lines for online programs specify

- the transaction
- the module name
- the control section name
- the DML command return offset
- the DML sequence number
- the function
- the version number

Note: To include a DML sequence number, the program must be compiled with the DEBUG DML compiler option.

The VIA column is not used in these report.

Figure 3-8. Attribution of CPU Execution Time Report

```
                                        ** ATTRIBUTION OF CPU EXECUTION TIME **

.IDMS    IDMSDBMS           IDMS DB MANAGER
--------------------------WAS INVOKED BY---------------------   -----------------VIA--------------------  CPU TIME %
XACTION  DIALOG      LOCATION    PROCESS NAME                     VERSION                                 SOLO  TOTAL

ADS2     PZTD2550        3200*   PZTD2550-PM                      0001                                     .07   .07
ADS2     PZTD2550       18000*   PZTD2550-PM                      0001                                     .00   .07
ADS2     PZTD2550       38500*   PZTD2550-PM                      0001                                     .00   .07
ADS2     PZTD2550       53800*   PZTD2550-PM                      0001                                     .39   .39
ADS2     PZTD2550       54100*   PZTD2550-PM                      0001                                     .34   .34
ADS2     PZTD2550        2000*1  PZTD2550-RP-ENTER                0001                                     .07   .07
ADS2     PZTD3492    00001260    PZTD3492-PM                      0001                                     .07   .07

XACTION  PROGRAM     SEQ#  FUNCTION
PAYROL2  PAYROL2     23    FIND NEXT SET                                                                  1.01   1.01
PAYROL2  PAYROL2     24    OBTN NEXT REC SET                                                               .85    .85
PAYROL2  PAYROL2     27    OBTN OWNER                                                                      .73    .73
.
XACTION  MODULE   SECTION  RETURN     SEQ#      FUNCTION                    VERSION
PAYROL3  PAYROL3           00000C5E   03        OBTN NEXT REC SET           00001                          1.01   1.01
                                                                                                         -----  -----
                                                                                                          4.68   4.68

.IDMS    IDMSINTD           IDMS/DBMS INTERFACE
--------------------------WAS INVOKED BY---------------------   -----------------VIA--------------------  CPU TIME %
XACTION  DIALOG      LOCATION    PROCESS NAME                     VERSION                                 SOLO  TOTAL

ADS2     PZTD2550       29900*   PZTD2550-PM                      0001                                     .07   .07
                                                                                                         ----- -----
                                                                                                           .07   .07


.IDMS    IDMSKEEP           IDMS DB KEY LOCKING RTN
--------------------------WAS INVOKED BY---------------------   -----------------VIA--------------------  CPU TIME %
XACTION  DIALOG      LOCATION    PROCESS NAME                     VERSION                                 SOLO  TOTAL

ADS2     PZTD2550       18000*   PZTD2550-PM                      0001                                     .07   .07
ADS2     PZTD2550       53800*   PZTD2550-PM                      0001                                     .07   .07
ADS2     PZTD2550       54100*   PZTD2550-PM                      0001                                     .14   .14
                                                                                                         ----- -----
                                                                                                           .29   .29


.IDMS    RHDCEVAL           RUNTIME EXP EVALUATION
--------------------------WAS INVOKED BY---------------------   --------------VIA-----------------------  CPU TIME %
XACTION  DIALOG      LOCATION    PROCESS NAME                     VERSION                                 SOLO  TOTAL

ADS2     PZTD1400-1  00000AF8    PZTD1400-PM                      0001                                     .07   .07
ADS2     PZTD2550        3200*   PZTD2550-PM                      0001                                     .14   .14
ADS2     PZTD2550       20500*   PZTD2550-PM                      0001                                     .07   .07
ADS2     PZTD2550       19800*   PZTD2550-PM                      0001                                     .35   .35
ADS2     PZTD2550       23000*   PZTD2550-PM                      0001                                     .36   .36
ADS2     PZTD2550       35000*   PZTD2550-PM                      0001                                     .07   .07
ADS2     PZTD2550       53800*   PZTD2550-PM                      0001                                     .30   .30
ADS2     PZTD2550       54100*   PZTD2550-PM                      0001                                     .28   .28
                                                                                                         ----- -----
                                                                                                          1.64  1.64
```

Total Line

The total line shows the total time attributed to the ADS/O dialogs that invoked the module. Because STROBE cannot always identify an invoker of a service routine, it may be less than the time shown in the Program Usage by Procedure or Most Intensively Executed Procedures reports.

To locate where you can improve performance, you should

- obtain an ADSORPTS listing of the dialog that shows the dialog's source
- find the statement(s) that caused the activity
- determine if the statement is coded efficiently by checking for complex arithmetical expressions or unnecessary calls to subroutines, dialogs, or DB/DC application programs
- modify the process source, regenerate the dialog, and create a Performance Profile for the new dialog
- compare the Performance Profiles to note changes in CPU consumption

Transaction Reports

The Transaction reports show you which transactions were responsible for most of the CPU activity. These reports are helpful when

- you want to identify which transactions (tasks) are responsible for a majority of the CPU consumption
- you want to see which modules within those transactions contributed the most to that transaction's CPU usage

Transaction Summary

The Transaction Summary report (Figure 3-9 on page 3-12) shows the CPU usage among transactions that were active during a measurement session. For a typical CA-IDMS-DC/CV region, this report shows four types of transactions:

- user-mode CA-IDMS applications, such as those written by database application developers and those supplied with the CA-IDMS system and ADS/O (ACCMNT, BYE and ADS2 shown in Figure 3-9 on page 3-12 are examples of these applications)
- system service tasks, such as those created by the CA-IDMS system to render support services (*DBRC*, *FACTOT*, *HELOT*, *LDRVR*, *MSTR* shown in Figure 3-9 on page 3-12 are examples of these tasks)
- external run units, such as Central Version batch jobs that issue database requests to the CA-IDMS CV for execution
- other transactions, such as activity in CA-IDMS TCBs that cannot be associated with a CA-IDMS task (CA-IDMS, for CPU activity in CA-IDMS main task, .IDMTSS, for CPU activity in a security system task, .RHDCTCK, for CPU activity in the CA-IDMS ticker task shown in Figure 3-9 on page 3-12 are examples of these transactions)

In addition to naming the transaction and defining its function, the Transaction Summary report gives TRANSACTION COUNT and AVERAGE SERVICE TIME. The TRANSACTION COUNT is the total number of times the transaction actually executed during the measurement session. The AVERAGE SERVICE TIME is the average wall clock time between transaction start and transaction end. These measurements enable you to see variances in the execution of a transaction. The CPU TIME and AVERAGE SERVICE TIME reported is based on STROBE sampling and are not directly related to the TRANSACTION COUNT.

As shown in the report, ADS/O task ADS2 using 36.31% and CA-IDMS application task ACCTMNT using 21.11% account for more than half of the total CPU usage. The reports shown in the next section “Transaction Usage by Control Section Reports” on page 3-12 provide detailed accounts of the CPU activity in these tasks.

Figure 3-9. Transaction Summary Report

```
                                                    ** TRANSACTION SUMMARY **

TRANSACTION          FUNCTION              TRANSACTION   AVERAGE      % CPU TIME     CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
    NAME                                      COUNT      SERVICE     SOLO   TOTAL   .00     10.00     20.00    30.00     40.00
                                                         TIME
  .IDMS          IDMS SYSTEM SERVICES                                5.29    6.87     .******++
  .IDMTSS        TOP SECRET                                          1.15    1.22     .*
  .RHDCTCK       TICKER TASK                                         7.68    7.97     .*******+
  *DBRC*         DATABASE RECOVERY TASK                              2.14    2.14     .**
  *ERUS*         EXT RUN-UNIT TASK                                   5.08    5.08     .*****
  *LDRVR*        LINE DRIVER TASK                                    4.22    4.29     .****
  *FACTOT*       FACTOTUM TASK                                       2.14    2.35     .**
  *HELOT*        HELOT TASK                                          1.01    1.11     .*
  *MSTR*         IDMS DC MASTER TASK                                 3.89    3.90     .****
  *PDRVR*        PRINTER LINE DRIVER TASK                             .07     .07     .
  ADS2           ADS RUNTIME PSEUDO-CONV          130     0.72    32.23   36.31     .**********************************++++
  ACCTMNT        ACCOUNT MAINTENANCE                             19.18   21.11     .*******************++
  ADS            ADS RUNTIME SYSTEM TASK            6     0.18     .07     .07     .
  PAYROL01       PAYROLL INPUT JOB                  2     0.44    3.73    3.77     .****
  PAYROL02       PAYROLL UPDATE JOB                 3     0.02    2.59    2.59     .****
  PAYROL03       PAYROLL ONLINE INQUERY             3     0.50    1.01    1.01     .*
  SIGNON         IDMS DC SIGNON TASK                              .14     .14     .
                                                --------    -----   -----   -----
PROGRAM IDMSPROD      TOTALS                       144     0.68   91.62  100.00
```

Note: The CA-IDMS Feature reporting of TRANSACTION COUNTS and AVERAGE SERVICE TIMES differs with the release of CA-IDMS that is active.

For CA-IDMS 10.2, the TRANSACTION COUNTS and AVERAGE SERVICE TIMES will be reported under the STROBE default pseudo-transaction name ".IDMS".

For CA-IDMS release 12.0, 12.01, and 14.0, the TRANSACTION COUNTS and AVERAGE SERVICE TIMES will be reported under the CA-IDMS supplied ERUS task named RHDCNP3S. The CA-IDMS 12.0, 12.01, and 14.0 architecture also supports use of a user-defined batch program name as a task to invoke the ERUS TASK program, which is also named RHDCNP3S. For a user-defined ERUS task, the TRANSACTION COUNTS and AVERAGE SERVICE TIMES will be reported under the user-defined task name. To use a user-defined batch program name, use either a SYSGEN or a DCMT commandTr.

Transaction Usage by Control Section Reports

The Transaction Usage by Control Section reports show, for each transaction that was active during the measurement session, the CPU time spent in each control section within each active module that was invoked to process that transaction.

Figure 3-10 on page 3-13 shows the CPU activity within the ADS/O task ADS2, the transaction under which ADS/O dialogs execute. This report details all application and system modules that executed on behalf of all ADS/O dialogs observed during the measurement session. It also includes any 3GL subroutines called by the dialogs such as ACCTINT, which is a CA-IDMS application module written in assembler that is called to calculate interest.

Figure 3-10. Transaction Usage by Control Section Report for ADS/O Tasks

```
                                    ** TRANSACTION USAGE BY CONTROL SECTION **


TRANSACTION ADS2              ADS RUNTIME PSEUDO-CONV

  MODULE      SECTION  COMPRESSED    FUNCTION                CPU TIME PERCENT    CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
   NAME         NAME                                             SOLO   TOTAL   .00      2.00      4.00      6.00      8.00

 .IDMS       IDMSDBIO  IDMSDBIO  IDMS DB I/O DRIVER             .29       .29     .*
 .IDMS       IDMSDBMS  IDMSDBMS  IDMS DB MANAGER               3.02      3.23     .***************+
 .IDMS       IDMSINTD  IDMSINTD  IDMS/DBMS INTERFACE            .07       .07     .
 .IDMS       IDMSKEEP  IDMSKEEP  IDMS DB KEY LOCKING RTN        .43       .43     .**
 .IDMS       IDMSPROD  IDMSPROD  IDMS SYSTEM SERVICES           .86       .93     .****
 .IDMS       RHDCAEDT  RHDCAEDT  OLM RUNTIME AUTO EDIT          .07       .07     .
 .IDMS       RHDCCSA             IDMS COMMON SYSTEM AREA       1.72      1.72     .********
 .IDMS       RHDCDBCO  RHDCDBCO  DB COORDINATION ROUTINE        .00       .07     .
 .IDMS       RHDCEDQ   RHDCEDQ   RESOURCE ENQUEUE ROUTINE       .29       .29     .*
 .IDMS       RHDCEVAL  RHDCEVAL  RUNTIME EXP EVALUATION        1.44      1.44     .*****
 .IDMS       RHDCLOCK  RHDCLOCK  LOCK/ILE MANAGER               .07       .07     .
 .IDMS       RHDCLODR  RHDCLODR  PROGRAM LOAD MANAGER          1.29      1.36     .*****
 .IDMS       RHDCMAPR  RHDCMAPR  OLM RUNTIME SERVICES          1.65      1.65     .******
 .IDMS       RHDCMISC  RHDCMISC  IDMS DC MISC SERVICES          .79       .79     .***
 .IDMS       RHDCPCTL  RHDCPCTL  PROGRAM CONTROL SERVICES       .07       .07     .
 .IDMS       RHDCRMGR  RHDCRMGR  RESOURCE CHAIN MGR              .43       .50     .**
 .IDMS       RHDCSCTY  RHDCSCTY  SECURITY/SIGNON PROCESS        .14       .14     .
 .IDMS       RHDCSTGP  RHDCSTGP  IDMS DC STORAGE MANAGER       1.87      2.16     .*********++++++
 .IDMS       RHDCTIMP  RHDCTIMP  IDMS DC TIME MGMT PGM          .07       .07     .
 .IDMS       RHDCTIRH  RHDCTIRH  BASIC MODE TERMINAL I/O        .07       .07     .
 .IDMS       RHDCWAIT  RHDCWAIT  MULTI-WAIT DISPATCHER         2.44      2.51     .****************+
 .NUCLEUS    IEAVRT04            CPU TIMER SERVICE             1.08      1.15     .*****+
 .NUCLEUS    IEAVTSFR            SETFRR SERVICE                 .07       .07     .
 .NUCLEUS    IEAVXSTK            S/B PCLINK STACK/UNSTACK       .93       .93     .*****
 .SVC        SVC 011             TIME                           .07       .14     .
 ACCTINT     ACCTINT             INTEREST CALCULATION RTN       .29       .29     .*
 ADSOMAIN    ADSOCCTL            RUNTIME CNTL CMD INTERP        .29       .36     .*
 ADSOMAIN    ADSOCDAB            RUNTIME DB CMD INTERP         1.51      1.51     .*******
 ADSOMAIN    ADSOCDRV            RUNTIME COMMAND DRIVER         .93      1.15     .****+
 ADSOMAIN    ADSOCEXC            RUNTIME EXEC CODE MGR          .36       .36     .*
 ADSOMAIN    ADSOCMAP            RUNTIME OLM CMD INTERP         .07       .07     .
 ADSOMAIN    ADSOCPRO            RUNTIME PROCEDURAL CMDS       2.23      2.37     .************+
 ADSOMAIN    ADSOCSAI            RUNTIME CSA INTERFACE          .07       .07     .
 ADSOMAIN    ADSOFPRO            RUNTIME CMD FAST PROCESS       .65       .72     .***
 ADSOMAIN    ADSOPTCH            RUNTIME PATCH MODULE           .14       .14     .
 ADSOMAIN    ADSORUN2            RUNTIME PSEUDO-CONVERSER      2.95      2.95     .****************
 ADSOMAIN    ADSORUN3            RUNTIME COMMON ROUTINES       2.51      2.59     .************
 ADSORUN1    ADSORUN1            RUNTIME ENVIRONMNT SETUP       .07       .14     .
 PMVECTRS    PMVECTRS            PERFORMANCE MONITOR            .93       .93     .****
                                                              -----     -----
TRANSACTION ADS2      TOTALS                                  32.23     36.31
```

Figure 3-11 on page 3-14 details the CPU activity used by all application and system modules that executed on behalf of the application task ACCTMNT. This task was created by the user and defined in the CA-IDMS Task Definition Table (TDT). Most of its activity is in the CA-IDMS database service modules IDMSDBMS, IDMSKEEP, and IDMSDBIO, indicating that this task is database intensive. ACCTMNT also invokes two user-written CA-IDMS programs called ACCTM100 and ACCTM200. Since both of these programs request DB/DC services primarily, the CPU activity directly attributed to them is relatively low.

Figure 3-11. Transaction Usage by Control Section Report for CA-IDMS Application Tasks

```
                                    ** TRANSACTION USAGE BY CONTROL SECTION **


TRANSACTION ACCTMNT

MODULE        SECTION  COMPRESSED     FUNCTION              CPU TIME PERCENT     CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
   NAME         NAME                                       SOLO     TOTAL        .00      2.00      4.00      6.00      8.00

 .IDMS         IDMSDBIO  IDMSDBIO   IDMS DB I/O DRIVER            2.51      2.73        .*************+
 .IDMS         IDMSDBMS  IDMSDBMS   IDMS DB MANAGER               6.18      6.97        .*************************++++
 .IDMS         IDMSINTD  IDMSINTD   IDMS/DBMS INTERFACE            .57       .72        .**+
 .IDMS         IDMSKEEP  IDMSKEEP   IDMS DB KEY LOCKING RTN       2.95      3.16        .**************+
 .IDMS         IDMSPROD  IDMSPROD   IDMS SYSTEM SERVICES           .22       .22        .*
 .IDMS         IDMSRDYX  IDMSRDYX   READY/READY EXTENT RTN         .29       .29        .*
 .IDMS         RHDCCSA              IDMS COMMON SYSTEM AREA       1.08      1.15        .*****
 .IDMS         RHDCCURS  RHDCCURS   CURRENCY SAVE/RESTORE          .79       .86        .***+
 .IDMS         RHDCEDQ   RHDCEDQ    RESOURCE ENQUEUE ROUTINE       .07       .07        .
 .IDMS         RHDCEVAL  RHDCEVAL   RUNTIME EXP EVALUATION         .14       .14        .
 .IDMS         RHDCLIMT  RHDCLIMT   RESOURCE LIMIT CHECK RTN       .14       .14        .
 .IDMS         RHDCLODR  RHDCLODR   PROGRAM LOAD MANAGER           .14       .14        .
 .IDMS         RHDCMISC  RHDCMISC   IDMS DC MISC SERVICES          .43       .43        .**
 .IDMS         RHDCRMGR  RHDCRMGR   RESOURCE CHAIN MGR             .36       .36        .*
 .IDMS         RHDCSCTY  RHDCSCTY   SECURITY/SIGNON PROCESS        .00       .07        .
 .IDMS         RHDCSTGP  RHDCSTGP   IDMS DC STORAGE MANAGER       1.01      1.22        .*****+
 .IDMS         RHDCWAIT  RHDCWAIT   MULTI-WAIT DISPATCHER         1.08      1.08        .*****
 .NUCLEUS      IEAVRT04             CPU TIMER SERVICE              .29       .43        .*+
 .NUCLEUS      IEAVTSFR             SETFRR SERVICE                 .14       .14        .
 .NUCLEUS      IEAVXSTK             S/B PCLINK STACK/UNSTACK       .07       .07        .
 ACCTM100                           ACCOUNT MAINTENANCE            .36       .36        .*
 ACCTM200                           ACCOUNT MAINTENANCE            .36       .36        .*
                                                                  -----     -----
TRANSACTION ACCTMNT   TOTALS                                      19.18     21.11
```

If programs show high CPU usage in the Transaction Usage by Control Section reports, you can see which sections of code within those programs cause the most CPU activity by referring to the appropriate Program Usage by Procedure reports.

Figure 3-12 on page 3-14 shows the CPU activity within the CA-IDMS-DC/CV region on behalf of the Central Version batch job PAYROL01. In this example, STROBE uses the batch job name as the transaction name and shows that this job invoked only CA-IDMS service modules. Since this job makes database requests to the CA-IDMS-DC/CV region, it mainly invokes CA-IDMS modules that provide database services. These batch jobs often access the same database as online tasks and can degrade performance for other online applications and CV batch jobs. Slow response time in the online applications and longer-running CV batch jobs may indicate that the same database areas and data sets are involved.

The STROBE Advanced Session Management Feature can assist in evaluating overall system performance. It enables you to measure jobs that are executing in different, but related address spaces. For example, if you use the STROBE Advanced Session Management Feature to measure the CV batch job and the CA-IDMS-DC/CV region simultaneously, you can compare their Performance Profiles and see how each of them affects the other. Using STROBE to measure a batch region is described in Chapter 4, “The STROBE Performance Profile for Batch Regions”.

Figure 3-12. Transaction Usage by Control Section Report for Central Version Batch Jobs

```
                                    ** TRANSACTION USAGE BY CONTROL SECTION **


TRANSACTION PAYROL01

  MODULE     SECTION  COMPRESSED    FUNCTION        CPU TIME PERCENT       CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
   NAME      NAME                                    SOLO     TOTAL        .00      2.00      4.00      6.00      8.00

 .IDMS     IDMSDBIO  IDMSDBIO      IDMS DB I/O DRIVER        1.02      1.02        .*****
 .IDMS     IDMSDBMS  IDMSDBMS      IDMS DB MANAGER           1.18      1.28        .*****+
 .IDMS     IDMSKEEP  IDMSKEEP      IDMS DB KEY LOCKING RTN   1.06      1.06        .*****
 .IDMS     RHDCCSA                 IDMS COMMON SYSTEM AREA    .47       .47        .**
                                                            -----     -----
TRANSACTION PAYROL01  TOTALS                                3.73      3.73
```

Figure 3-13 details the CPU activity used by the CA-IDMS FACTOTUM task *FACTOT*. This task was created by CA-IDMS to handle system-level services such as terminal paging, as indicated by the asterisks STROBE placed around the pseudo-taskcode name. The report reflects all CPU activity caused by FACTOTUM tasks observed in the CA-IDMS system during the measurement session. Detailing these system service pseudo-tasks shows the impact that they have on CA-IDMS system modules. These service tasks are embedded in the CA-IDMS system and support many functions used by applications, but they cannot be modified to improve Tperformance.

Figure 3-13. Transaction Usage by Control Section Report for CA-IDMS Tasks

```
                                    ** TRANSACTION USAGE BY CONTROL SECTION **


TRANSACTION *FACTOT*          FACTOTUM TASK

  MODULE   SECTION  COMPRESSED  FUNCTION           CPU TIME PERCENT       CPU TIME HISTOGRAM  MARGIN OF ERROR:  2.63%
   NAME    NAME                                     SOLO     TOTAL        .00      2.00      4.00      6.00      8.00

 .IDMS    RHDCCSA               IDMS COMMON SYSTEM AREA      .14       .21        .*
 .IDMS    RHDCMSTR  RHDCMSTR    MASTER-TASK CNTL PRGM        .07       .07        .
 .IDMS    RHDCRMGR  RHDCRMGR    RESOURCE CHAIN MGR           .14       .14        .*
 .IDMS    RHDCSCRM  RHDCSCRM    SCRATCH MANAGER              .14       .14        .*
 .IDMS    RHDCSTGP  RHDCSTGP    IDMS DC STORAGE MANAGER      .14       .14        .*
 .IDMS    RHDCWAIT  RHDCWAIT    MULTI-WAIT DISPATCHER       1.51      1.65        .******+
                                                            -----     -----
TRANSACTION *FACTOT*  TOTALS                                2.14      2.35
```

Wait Time by Module Report

In a CA-IDMS-DC/CV region, wait time is shown in the CA-IDMS main task, RHDCTCKR, and the RHDCCKUR task. The CA-IDMS main task will wait for operating system or subtask services, such as program loading, to complete. The STROBE CA-IDMS Feature identifies, in the Attribution of CPU Wait Time report, the module and control section responsible for causing this wait.

The module function descriptor is a good indicator of the service that caused the wait. For example, wait detected in control section BLDLLOAD is generated by a large number of program loads or an unusually high number of concatenated LOADLIB data sets. You can investigate program residency and also try to reduce the number of concatenated LOADLIBs to reduce wait time associated with program loads.

Both the RHDCTCKR and the RHDCCKUR tasks execute CA-IDMS-supplied code and manage runaway task detection and external run-unit activity, respectively. They spend the majority of time waiting for system-oriented events to occur and present no real performance improvement opportunities.

The Performance Profile shows wait time in CA-IDMS address spaces in two reports, the Wait Time by Module report and the Attribution of CPU Wait Time report.

The Wait Time by Module report (Figure 3-14 on page 3-16) shows the sites of wait. In this example, all wait time is attributed to the CA-IDMS module RHDCCKUR, which waits while an external run unit is active.

Figure 3-14. Wait Time by Module Report

```
                                        ** WAIT TIME BY MODULE **

 MODULE   SECTION COMPRESSED    FUNCTION           RUN TIME PERCENT       RUN TIME HISTOGRAM  MARGIN OF ERROR;   .62%
  NAME      NAME    SECTION                        PAGE     TOTAL        .00     16.50     33.00     49.50     66.00

 .IDMS    RHDCCKUR            IDMS CHECK USER PGM TASK           1.86     64.44        .+++++++++++++++++++++++++++
                                                               -----     -----
PROGRAM IDMSPROD TOTALS                                         1.86     64.44
```

Reports That Show I/O Utilization

You can identify I/O performance improvement opportunities in CA-IDMS by interpreting STROBE reports according to your knowledge of the target database's organization — the relationship among areas, sets, records, and ddnames. By taking the database organization into account, you can use the STROBE reports to gain insight into the principal I/O loads operating against the database.

You can also determine how a particular database is accessed, which can lead to performance improvement opportunities. For example, high wait time associated with the CA-IDMS load area in a CA-IDMS-DC/CV region that is primarily ADS/O can indicate frequent loads of dialogs and their associated maps. By link editing the dialogs and maps into load libraries, you can reduce the contention for the load area and not adversely affect the performance of library loads by considering program residency for the more frequently used programs.

Resource Demand Distribution

The I/O reports show all the data sets active in the address space. The Resource Demand Distribution report (Figure 3-15) shows the wait time attributed to CA-IDMS data sets. From this example, you can see that the data definition name DLODDB (the CA-IDMS load area data set) is responsible for most of the wait attributed to the CA-IDMS data sets. You can then use the I/O Facility Utilization Summary report to find more details about ddname DLODDB.

Figure 3-15. Resource Demand Distribution Report

```
                                        ** RESOURCE DEMAND DISTRIBUTION **

                      ---- PERCENT OF RUN TIME -----    ----- PERCENT OF RUN TIME SPENT ----   CUMULATIVE PERCENTAGES
TASK OR   RESOURCE      SERVICED  SERVICED  SERVICED      SOLO     SOLO      SOLO     CAUSING      SOLO     CAUSING
 DDNAME                  BY CPU    BY I/O  BY EITHER     IN CPU   IN I/O  IN EITHER   CPU WAIT     TIME    CPU WAIT

IDMSPROD  CPU              5.09       .00      5.09       4.63      .00      4.63        .00      4.63        .00
RHDCTCKR  CPU               .39       .00       .39        .37      .00       .37        .00      5.00        .00
IDMTSS    CPU               .07       .00       .07        .06      .00       .06        .00      5.06        .00
RHDCCKUR  CPU               .00       .00       .00        .00      .00       .00      76.70      5.06      76.70
RHDCCKUR1 CPU               .00       .00       .00        .00      .00       .00      10.48      5.06      87.18
DLODDB    3380              .00      2.95      2.95        .00     2.48      2.48       2.74      7.54      89.92
PZTORG    3390              .00      1.08      1.08        .00      .94       .94        .94      8.48      90.86
PZTCLM    3390              .00       .99       .99        .00      .93       .93        .93      9.41      91.79
R500      3390              .00       .63       .63        .00      .56       .56        .56      9.97      92.35
DLOGDB    3380              .00       .62       .62        .00      .47       .47        .50     10.44      92.85
DICTDB    3380              .00       .48       .48        .00      .46       .46        .46     10.90      93.31
DSCDB     3380              .00       .23       .23        .00      .22       .22        .22     11.12      93.53
PZTTABLE  3390              .00       .25       .25        .00      .16       .16        .16     11.28      93.69
PZTPCLM   3390              .00       .10       .10        .00      .09       .09        .09     11.37      93.78
J1JRNL    3380              .00       .07       .07        .00      .07       .07        .07     11.44      93.85
  OTHER   I/O               .02       .66       .68        .02      .58       .60        .58     12.04      94.43
```

I/O Facility Utilization

The I/O Facility Utilization Summary report (Figure 3-16) shows I/O activity by device, volume, and ddname. The DASD Usage by Cylinder report details DASD usage and shows I/O activity by device, volume, ddname, and cylinder.

Figure 3-16. I/O Facility Utilization Summary Report

```
                                    ** I/O FACILITY UTILIZATION SUMMARY **

 UNIT      DEVICE  CACHE  VOLUME     DDNAME           RUN TIME PERCENT       RUN TIME HISTOGRAM  MARGIN OF ERROR:   .62%
  NO         TYPE   ELIG   ID                           SOLO     TOTAL        .00      1.00      2.00      3.00      4.00

 C56     DA 3380          DBMS10    STEPLIB            .01       .01        .

 214     DA 3380          MVS603    DICTDB             .46       .48        .****
 214     DA 3380          MVS603    DLODDB            2.48      2.95        .*************************+++++
 214     DA 3380          MVS603    DLOGDB             .47       .62        .****++
 214     DA 3380          MVS603    DQUDB              .01       .01        .
 214     DA 3380          MVS603    DSCDB              .22       .23        .**
                                                     -----     -----
UNIT  214 TOTALS                                     3.64      4.29

 600     DA 3380          MVS103    CDMSLIB            .04       .04        .

 603     DA 3380          MVS502    J1JRNL             .07       .07        .

 607     DA 3380          MVS700    DMSGDB             .04       .04        .

 73A     DA 3390          DPAC24    PZTIXDOC           .02       .02        .

 73B     DA 3390          RLOD22    PZTACTV            .04       .04        .
 73B     DA 3390          RLOD22    PZTAPOL            .02       .02        .
 73B     DA 3390          RLOD22    PZTCLM             .93       .99        .*********
 73B     DA 3390          RLOD22    PZTGLBL            .00       .00        .
 73B     DA 3390          RLOD22    PZTINJ             .03       .04        .
 73B     DA 3390          RLOD22    PZTIXCNT           .05       .05        .
 73B     DA 3390          RLOD22    PZTIXDRY           .03       .03        .
 73B     DA 3390          RLOD22    PZTIXFSG           .00       .01        .
 73B     DA 3390          RLOD22    PZTIXLOC           .01       .02        .
 73B     DA 3390          RLOD22    PZTIXNAM           .02       .02        .
 73B     DA 3390          RLOD22    PZTLNDR            .02       .02        .
 73B     DA 3390          RLOD22    PZTNOTE            .03       .03        .
 73B     DA 3390          RLOD22    PZTNTXT3           .02       .02        .
 73B     DA 3390          RLOD22    PZTORG             .94      1.08        .*********+
 73B     DA 3390          RLOD22    PZTPCLM            .09       .10        .+
 73B     DA 3390          RLOD22    PZTPROD            .04       .04        .
 73B     DA 3390          RLOD22    PZTPROP            .05       .06        .
 73B     DA 3390          RLOD22    PZTPRPT            .05       .05        .
 73B     DA 3390          RLOD22    PZTREP             .00       .00        .
 73B     DA 3390          RLOD22    PZTSADM            .03       .03        .
 73B     DA 3390          RLOD22    PZTTABLE           .16       .25        .*+
 73B     DA 3390          RLOD22    PZTTRAN            .02       .03        .
 73B     DA 3390          RLOD22    PZTWRKL            .00       .01        .
 73B     DA 3390          RLOD22    PZTZIP             .00       .01        .
 73B     DA 3390          RLOD22    R500               .56       .63        .*****+
                                                     -----     -----
UNIT  7 B TOTALS                                     3.14      3.58

 73E     DA 3390          DPAC23    PZTCASH            .00       .01        .
```

The data sets PZTCLM, PZTORG, and R500 are located on one volume and show the highest activity. Consider relocating them to different volumes with less activity to balance the I/O load. Additionally, the application developer's knowledge can help in locating and resolving performance problems. For example, dialog PZTD2550, process PZTD2550-PM (Figure 3-5 on page 3-7) shows statements 18000, 18300, and 20500 issuing DML commands against sets that are within the data set PZTCLM. Process PZTD2550-RP-ENTER statement 5000 also accesses PZTCLM. You can consider using a database analyzer utility to determine the organization of sets and records within PZTCLM and whether area reorganization may improve database I/O performance.

Data Set Characteristics

The Data Set Characteristics report (Figure 3-17) shows, for each data set accessed during the measurement session, the access method, the record format and size, block size, number of buffers, and data set name. The number of I/O operations performed (EXCPs)

during the measurement session appears as well, showing which data sets are most often accessed. You can use it as a starting point for evaluation of database efficiency.

Figure 3-17. Data Set Characteristics Report

```
                                        ** DATA SET CHARACTERISTICS **

DDNAME     ACCESS     POOL    REC   BLK/CI   HBUF  BUF    RPL  - SPLITS-      EXCP  DATA SET NAME
           METHOD       NO    SIZE   SIZE      NO   NO STRINGS  CI    CA    COUNTS

CDMSLIB  BPAM                      32760                                            IDMS.PROD.GENLOAD
CDMSLIB  BPAM                      32760                                            IDMS.PROD.LOAD
CDMSLIB  BPAM                      32760                                            IDMS.DLB.LOAD
CDMSLIB  BPAM                      32760                                            IDMS.PROD.NDVRLOAD
CDMSLIB  BPAM                      32760                                            IDMS.DCT1.DBA5LOAD
CDMSLIB  BPAM                      32760                                        7   SYS1.IDMS.COB2LIB
CDMSLIB  BPAM                      32760                                            SYS1.TSS.LINKLIB
PZTACTV  BDAM                       3860                                       11   IDMS.PROD.PZTACTV
PZTAPOL  BDAM                       7476                                        7   IDMS.PROD.PZTAPOL
PZTCASH  BDAM                       3860                                        4   IDMS.PROD.PZTCASH
PZTCLM   BDAM                       7476                                      252   IDMS.PROD.PZTCLM
PZTGLBL  BDAM                       3860                                        1   IDMS.PROD.PZTGLBL
PZTINJ   BDAM                       3188                                        8   IDMS.PROD.PZTINJ
PZTIXCNT BDAM                       7476                                       12   IDMS.PROD.PZTIXCNT
PZTIXDOC BDAM                       2676                                        2   IDMS.PROD.PZTIXDOC
PZTIXDRY BDAM                       3860                                        7   IDMS.PROD.PZTIXDRY
PZTIXFSG BDAM                       7476                                        2   IDMS.PROD.PZTIXFSG
PZTIXLOC BDAM                       7476                                        4   IDMS.PROD.PZTIXLOC
PZTIXNAM BDAM                       7476                                       10   IDMS.PROD.PZTIXNAM
PZTLNDR  BDAM                      23476                                        2   IDMS.PROD.PZTLNDR
PZTNOTE  BDAM                      23476                                        7   IDMS.PROD.PZTNOTE
PZTNTXT3 BDAM                       7476                                        8   IDMS.PROD.PZTNTXT3
PZTORG   BDAM                       3860                                      489   IDMS.PROD.PZTORG
PZTPCLM  BDAM                       3188                                       33   IDMS.PROD.PZTPCLM
PZTPROD  BDAM                       7476                                        7   IDMS.PROD.PZTPROD
PZTPROP  BDAM                       7476                                       17   IDMS.PROD.PZTPROP
PZTPRPT  BDAM                       3188                                       13   IDMS.PROD.PZTPRPT
PZTREP   VSAM ESDS           23476  24576            4     3                       IDMS.PROD.PZTREP
PZTSADM  BDAM                       7476                                        5   IDMS.PROD.PZTSADM
PZTTABLE BDAM                      23476                                       80   IDMS.PROD.PZTTABLE
PZTTRAN  BDAM                       3860                                        6   IDMS.PROD.PZTTRAN
PZTWRKL  BDAM                       7476                                        4   IDMS.PROD.PZTWRKL
PZTZIP   BDAM                       2676                                        2   IDMS.PROD.PZTZIP
DICTDB   BDAM                       7476                                       28   IDMS.PROD.DDLDML
DLODDB   BDAM                      11476                                      378   IDMS.PROD.DDLDCLOD
DLOGDB   BDAM                       3860                                       62   IDMS.PROD.DDLDCLOG
DMSGDB   BDAM                       3860                                       15   IDMS.DCP6.DDLDPZTG
DQUDB    BDAM                       2676                                        1   IDMS.PROD.DDLDCRUN
DSCDB    BDAM                       2676                                       43   IDMS.PROD.DDLDCSCR
R500     BDAM                      11476                                      139   IDMS.DCS5.R500
J1JRNL   BDAM                       2004                                       88   IDMS.PROD.JOURNAL1
STEPLIB  BPAM                      23476                                            IDMS.PROD.GENLOAD
```

Chapter 4.
The STROBE Performance Profile for Batch Regions

CA-IDMS Central Version and local mode batch regions are characterized by high amounts of wait time because these jobs tend to be I/O-intensive, with nearly all of the I/O resulting from CA-IDMS DB DML commands issued by an application program. The STROBE CA-IDMS Feature assigns the majority of this wait time to IDMSINTB, the CA-IDMS Batch Interface module, and provides a more detailed breakdown by DML command. STROBE identifies, by program name, control section name, offset within the program, and, if available, DML sequence number, each DML command responsible for causing wait.

This chapter provides three Performance Profile reports generated for a batch job running CA-IDMS. The report data applies to releases 10.2, 12.0, 12.01, and 14.0 of CA-IDMS.

Measurement Session Data Report

The Measurement Session Data report (Figure 4-1) describes the environment during a measurement session. The low CPS time percent in this example is typical of Central Version batch mode because all CA-IDMS database activity is performed by the CA-IDMS-CV address space. In the SUBSYSTEM field, the name IDMS is followed by its release number, which is followed by a three-character code indicating whether the address space is running Central Version batch (code “CVB”) or local mode (code “LOC”). If no characters appear, STROBE measured an online region.

Figure 4-1. Measurement Session Data Report

```
                                        ** MEASUREMENT SESSION DATA **

------- JOB ENVIRONMENT --------        ------- MEASUREMENT PARAMETERS ------        ----- MEASUREMENT STATISTICS ----

PROGRAM MEASURED   -    IDMSB757       ESTIMATED SESSION TIME  -      40 MIN        CPS TIME PERCENT  -          1.29
JOB NAME           -    SANCP550       TARGET SAMPLE SIZE      -      50,000        WAIT TIME PERCENT -         98.71
JOB NUMBER         -    JOB04578       REQUEST NUMBER          -         739        RUN MARGIN OF ERROR PCT -     .67
STEP NAME   -  PPM50550.ID757S25                                                    CPU MARGIN OF ERROR PCT -     .15

                                       MODULE MAPPING BASELINE -           0        TOTAL SAMPLES TAKEN -      21,193
DATE OF SESSION    -    06/23/96       ATTRIBUTORS    -              ATTRSVC        TOTAL SAMPLES PROCESSED -  21,193
TIME OF SESSION    -    10.24.10                                                    INITIAL SAMPLING RATE - 20.83/SEC
CONDITION CODE     -      C-0000                                                    FINAL SAMPLING RATE   - 20.83/SEC
                                       --------- REPORT PARAMETERS ---------
SYSTEM -             ESA SP4.2.2                                                   SESSION TIME -   21 MIN  3.28 SEC
DFSMS -                    1.1.0       REPORT RESOLUTION    -       64 BYTES      CPU TIME -        0 MIN 23.09 SEC
SUBSYSTEM          - IDMS 10.2CVB      SORTSIZE             -       100,000       WAIT TIME      - 19 MIN 09.96 SEC
CPU MODEL          -   3090-600J       LINES/PAGE           -            60       STRETCH TIME -    1 MIN 30.23 SEC
SMF/SYSTEM ID      -         IAS6
                                       DASD=2.0% DASDGAP=5                         SRB TIME -       18 MIN 14.69 SEC
REGION SIZE BELOW 16M  -  4,160K       ATTR= 0.0%                                  SERVICE UNITS -               542
REGION SIZE ABOVE  -     32,768K
                                       DATE FORMAT                MM/DD/YYYY       PAGES IN -    71     OUT-       1
PTF LVL- 2.5.0.FS000000/FS000000       TIME FORMAT                  HH:MM:SS       PAGING RATE  -           0.06/SEC
STROBE TAPE NUMBER -  000-S85999                                                   EXCPS -      1,657       1.44/SEC

SAMPLE DATA SET  - STROBE.SAMPPROF.S003D001
```

Wait Time by Module

The Wait Time by Module report for Batch Regions (Figure 4-2) shows the sites of wait time. In this CA-IDMS-CV batch address space, almost all of the wait time is shown in the pseudo-module .IDMS as attributed to the control section IDMSINTB, the CA-IDMS Batch Interface. This wait time represents the time incurred by the CV region to service the database request.

Figure 4-2. Wait Time by Module Report for Batch Regions

```
                                        ** WAIT TIME BY MODULE **

 MODULE SECTION COMPRESSED    FUNCTION                 RUN TIME PERCENT       RUN TIME HISTOGRAM  MARGIN OF ERROR:  .67
  NAME      NAME  SECTION                              PAGE     TOTAL        .00     23.50     47.00     70.50    94.00

 .COBLIB  ILBOCMM           GETCORE                    .00       .00        .
 .COBLIB  ILBODSP           DISPLAY                    .00       .00        .
                                                       -----     -----
 .COBLIB  TOTALS            COBOL LIBRARY SUBROUTINE   .00       .00

 .IDMS    IDMSINTB          IDMS/BATCH INTERFACE       .11     92.50        .+++++++++++++++++++++++++++++++++++++

 .IOCS    IGG019AQ          QSAM GET, NEXT BUFFER      .00       .55        .
 .IOCS    IGG019AR          QSAM PUT, NEXT BUFFER      .00      1.99        .
 .IOCS    IGG019BB          BSAM CHECK ALL DEVICES     .00       .02        .
 .IOCS    IGG019LI          BDAM CHECK MODULE          .00       .18        .
                                                       -----     -----
 .IOCS    TOTALS            DATA MANAGEMENT SERVICES   .00      2.74

 .LKD/LDR IEWFETCH          PROGRAM FETCH              .00       .01        .

 .SVC     SVC 006           PROGRAM MANAGER/LINK       .00       .08        .
 .SVC     SVC 008           PROGRAM MANAGER/LOAD       .00       .15        .
 .SVC     SVC 018           BLDL/FIND                  .00      2.56        .+
 .SVC     SVC 019           OPEN                       .00       .03        .
 .SVC     SVC 020           CLOSE                      .03       .14        .
 .SVC     SVC 026           CATALOG MANAGEMENT         .00       .19        .
 .SVC     SVC 099           DYNAMIC ALLOCATION         .03       .08        .
 .SVC     SVC 111           JES 2                      .00       .15        .
                                                       -----     -----
 .SVC     TOTALS            SUPERVISOR CONTROL         .06      3.38

 .VSAM    IDA019L1          VSAM RECORD MANAGEMENT     .00       .01        .
 CPXCEM                     CA/OPTIMIZER ANALYZER      .00       .00        .
 CPXTERM                                               .00       .00        .
 ERBMFRES                                              .00       .00        .
 HASPSSSM                   SYSIN/SYSOUT DATA HANDLR   .01       .01        .
 IDMSB757 IDMSB757                                     .00       .01        .

 TBNUCL   GC0103                                       .01       .01        .
 TBNUCL   GC0432                                       .00       .00        .
                                                       -----     -----
 TBNUCL   TOTALS                                       .01       .01

 VIOCLX41                                              .00       .00        .
                                                       -----     -----
PROGRAM IDMSB757 TOTALS                                .19     98.67
```

Attribution of CPU Wait Time Report

The Attribution of CPU Wait Time report (Figure 4-3 on page 4-3) identifies the sites of invocation of system routines. When wait time is attributed to the CA-IDMS Batch Interface IDMSINTB invokers, the report header displays the invoking routine, showing

- its pseudo-module, module, and control section name (when available)
- a function descriptor for either the control section or the module

Each report detail line identifies an invoker of the batch wait and includes

- the module and control section name
- a DML command return offset

- a DML sequence number, if the program was compiled with the DEBUG DML compiler option
- page wait and total wait time caused by the invoker

Figure 4-3. Attribution of CPU Wait Time Report

```
                                        ** ATTRIBUTION OF CPU WAIT TIME **

.IDMS    IDMSINTB IDMSINTB  IDMS/BATCH INTERFACE
--------------------------WAS INVOKED BY--------------------   -----------------VIA----------------------- WAIT TIME %
XACTION  MODULE   SECTION    RETURN     SEQ# PROCEDURE NAME       MODULE   SECTION  FUNCTION                PAGE  TOTAL

         CDIOP007            002466       1                                                                  .00   .04
         CDIOP007            0025F2       3                                                                  .00   .08
         CDIOP007            002A26       4                                                                  .00   .01
         IDMSB757 IDMSB757   008DC4       6                                                                  .00   .02
         IDMSB757 IDMSB757   0093B6      11                                                                  .00   .01
         IDMSB757 IDMSB757   00A666      12                                                                  .02 26.40
         IDMSB757 IDMSB757   00A738      13                                                                  .01  7.28
         IDMSB757 IDMSB757   00A7AA      14                                                                  .00  2.16
         IDMSB757 IDMSB757   00A956      15                                                                  .07 53.52
         IDMSB757 IDMSB757   00A9C0      16                                                                  .00  2.94
         IDMSB757 IDMSB757   00D68A      17                                                                  .00   .04
                                                                                                            ----- -----
                                                                                                              .11 92.50
```

In this example, control section IDMSB757 in module IDMSB757 accounts for 92.37% of the wait time incurred by the CA-IDMS batch interface. Two areas within this control section are of particular interest: the call statement at offset 00A956, DML sequence number 15, accounts for 53.52% of the wait time; and the call statement at offset 00A666, DML sequence number 12, accounts for 26.40% of the wait time.

The offset or DML sequence number can be compared with the compiler listing to determine the line of code that is causing the wait. The sample code shown in Figure 4-4 shows DML sequence number 15, at offset 00A956. This offset falls between offsets 00A940 and 00A9B6 in the code, and it is actually the offset to which CA-IDMS returns control within the group of instructions generated by the COBOL compiler when it compiled the CALL verb. Therefore, when you compare the Attribution report to your listing, examine the group of statements indicated by the DML sequence number rather than the exact offset listed.

Figure 4-4. CA-IDMS/COBOL Code Showing the DML Sequence Numbers

```
 1060-FIND-VIA-SYMBOL                 SECTION.                    00034560 00A8EE
     MOVE '1060-FIND-VIA-SYMBOL' TO PM-ABEND-PARAGRAPH.           00034570 00A924
                                                                  00034580
                                                                  00034590
     MOVE TS-UNIQUE-SYMBOL TO UNIQUE-SYNONYM-KEY-1704.            00034600 00A934
#    OBTAIN  GS-DNASYN-1704  WITHIN  IX-PM-DNASYN                 00034610
#     USING UNIQUE-SYNONYM-KEY-1704.                              00034620
                MOVE 0015 TO DML-SEQUENCE                         DMLC0015 00A93A
                CALL 'IDMS' USING SUBSCHEMA-CTRL                          00A940
                                  IDBMSCOM (33)
                                  SR1704
                                  IX-PM-DNASYN
                                  UNIQUE-SYNONYM-KEY-1704
                                  IDBMSCOM (43).
                                                                  00034630
     IF DB-STATUS-OK                                              00034640 00A9B6
```

In this case, the statement is a record retrieval from an indexed or sorted set. The amount of wait time indicates that the set's structure should be investigated. You should also determine the number of concurrent access attempts.

Attribution of CPU Execution Time Report

The Attribution of CPU Execution Time report (Figure 4-5) identifies the sites of invocation of system routines. When execution time is attributed to the CA-IDMS Batch Interface IDMSDBMS invokers, the report header displays the invoking routine, showing

- its pseudo-module, module, and control section name (when available)
- a function descriptor for either the control section or the module

Each report detail line identifies an invoker causing CPU execution time and includes

- the module and control section name
- a DML command return offset
- a DML sequence number, if the program was compiled with the DEBUG DML compiler option
- execution time and total execution time caused by the invoker

Figure 4-5. Attribution of CPU Execution Time Report

```
                                        ** ATTRIBUTION OF CPU EXECUTION TIME **

.IDMS    IDMSDBMS IDMSDBMS  IDMS/BATCH INTERFACE
--------------------------WAS INVOKED BY---------------------   -----------------VIA----------------------- CPU TIME %
XACTION  MODULE   SECTION    RETURN     SEQ# PROCEDURE NAME       MODULE   SECTION  FUNCTION                SOLO  TOTAL

         E2ANDLEL E2ANDLEL   00102A       16                                                                  1.08  1.08
         E2ANDLEL E2ANDLEL   00115A       19                                                                   .72   .72
         E2ANDLEL E2ANDLEL   0011BE       20                                                                  1.08  1.08
         E2ANDLEL E2ANDLEL   001264       21                                                                  1.43  1.43
         E2ANDLEL E2ANDLEL   0012D2       22                                                                  1.43  1.43
         E2ANDLEL E2ANDLEL   00138E       23                                                                   .36   .36
         E2ANDLEL E2ANDLEL   0013FA       24                                                                  2.15  2.15
                                                                                                              ----------
                                                                                                              8.24  8.24
```

In this example, control section E2ANDLEL in module E2ANDLEL accounts for 8.24% of the execution time incurred by the CA-IDMS batch interface. The call statement at offset 0013FA, DML sequence number 24, accounted for the largest amount of execution time for the control section execution time using 2.15%.

Using Database Analysis Utilities with STROBE

To help identify the source of performance problems, you can combine information from the STROBE Performance Profile and your database utilities. As shown in the previous example, if the problem lies in record retrieval, then your database utilities can give you the average set length and the number of levels in your index. You can derive other helpful information by issuing the CA-IDMS command ACCEPT STATISTICS before your job has completed processing.

Sorted Sets

If the set is sorted, with many members chained off each set “owner”, and access by sort key is random, you can decrease the amount of search time by

- adding a set from the member record that contains an index to the member specifying the same symbolic key
- transforming the set from a sorted set into an indexed set
- reducing the number of members per owner

Indexed Sets

If the set is indexed, then the index may be fragmented or unbalanced, or there may be collisions between batch and online users who need to update the index. To determine whether several users are trying to access the index concurrently, submit simultaneous STROBE measurement sessions for both the batch program and the CV region. Then examine the overlap between the Performance Profiles to determine how many transactions are trying to access the index files or whether your database areas are on the same physical units as your database journals and log files.

In these cases, you can

- rebuild the index as sorted so that it is more evenly balanced
- ensure that all applications updating the index lock and release the specific records they need (COMMIT frequently) to allow other users concurrent access to other records
- check usage modes to ensure that all users browsing the index are not holding update locks
- ensure that your batch job issues frequent COMMIT statements that release any locked records
- consider duplicating index information so that online users can browse an updated index without conflicting with the operation of the batch job, then limit the online access so that it becomes complementary

CALC or VIA

If your sets are stored using a CALC key or VIA a set owner, database performance can degrade when the assigned database area becomes heavily populated (database statistics will show a high number of CALC and VIA overflows). In this case, you should unload the database, enlarge the database areas, and reload the database.

Database Buffering

If database statistics show a narrow ratio between pages-read and pages-requested, there may be insufficient database buffer pages or conflicting utilization of buffer pools. You can find this information by issuing the CA-IDMS command ACCEPT STATISTICS before your job has completed processing.

Index

Special Characters

DBRC, 1-5
ERUS, 1-5
FACTOT, 1-5, 3-15
HELOT, 1-5
LDRVR, 1-5
MSTR, 1-5
PDRVR, 1-5
.IDMS
  activity attributed to, 1-5
  in the Wait Time by Module report for batch regions, 4-2
  reporting of transaction counts and average service times for CA-IDMS 10.2, 3-12
.SYSTEM
  activity attributed to, 1-5
  in the Program Section Usage Summary report, 3-3

A

ACCEPT STATISTICS, IDMS command
  in determining sufficient database buffer pages or pools, 4-5
  use in identifying performance problems, 4-4
ADD command, invoking or suspending data collectors, 2-1–2-2
ADS/O
  described, 1-3
  dialogs
    diagnostic table requirement, 2-3, 3-7
    indexing, 2-2–2-5
    Performance Profile reports that indicate CPU activity, 1-3
    reported on by the STROBE CA-IDMS Feature, 1-3
    reports that show CPU activity in, 3-5
    statement numbers
      diagnostic tables requirement, 2-3, 3-7
  disabling data collector, 2-1–2-2
ADSORPTS utility program
  concatenate listings, 2-5
  parameters values, 2-3
  use of described, 2-3
ATTR, keyword parameter, 2-5
attribution data
  disabling CA-IDMS data collection, 2-1
  limiting collection of, 2-5
Attribution of CPU Execution Time report
  described, 3-2
  explained, 3-9–3-10
  figure of, 3-10
  suppressing, 2-5
Attribution of CPU Wait Time report
  described, 4-3–4-4
  suppressing, 2-5
  use in improving performance, 1-4
attribution reports, suppressing, 2-5–2-6

B

baseline, specifying for wait or CPU in attribution reports, 2-5
batch, CA-IDMS modes
  Central Version mode, 1-3
  indexing in, 2-4
  local mode, 1-4
buffer pages and buffer pools, adjusting, 4-5

C

CA-IDMS
  ADSORPTS utility program
    parameters for, 2-3
  Batch Central Version mode, described, 1-3
  Batch Interface (IDMSINTB)
    in the Attribution of CPU Wait Time report, 4-2
    in the Wait Time by Module report, 4-2
    wait time reported to, 4-1
  batch local mode, described, 1-4
  dialogs, ADS/O, 1-3
  environments, described, 1-2–1-4
  keyword parameters, 2-1
  measurement requests, submitting, 2-1–2-2
  programs, attribution of CPU execution time, 4-4
  programs, attribution of wait time, 4-2
  regions described, 1-1
CA-IDMS ADSORPTS utility program
  concatenate listings, 2-5
  use of described, 2-3
CA-IDMS-CV tasks, 1-3
CA-IDMS-DC tasks, 1-2–1-3
CA-IDMS-DC/CV mode, described, 1-2–1-3
CA-IDMS-DC/CV-UCF mode, described, 1-4
CALC, using as a set key, 4-5
CICS
  using with CA-IDMS-DC/CV-UCF, 1-4
CPU activity
  ADS/O dialogs responsible for, 1-3, 3-5
  ADS/O statements responsible for, 1-3, 3-6
  application programs responsible for, 3-4
  control sections responsible for, 1-3, 3-3, 3-12
  reported for application programs, 3-4
  summary of reports that show, 3-2
  system service routines responsible for, 3-4
  transactions responsible for, 3-11
CPU Usage by ADS/O Statement report
  described, 3-6–3-8
  use in improving performance, 1-3, 3-2

D

data collection
disabling, 2-1–2-2
of CA-IDMS data, 2-1–2-2
Data Set Characteristics report
explained, 3-17
figure of, 3-18
database service tasks, CA-IDMS-CV services, 1-2–1-3
DEBUG, DML compiler option requirement, 2-2, 3-9
DETAIL parameter, 1-5
diagnostic tables, required for statement numbers, 2-2, 3-7
Dialog CPU Summary report
explained, 3-4–3-6
identified, 3-2
use in performance improvement, 1-3
dialogs, ADS/O dialogs defined, 1-3
DIALOGS, ADSORPTS parameter, 2-3
DML
sequence numbers, 1-4, 4-3
DEBUG DML compiler option requirement, 3-9
statement offsets, 1-4

E

environments, for IDMS
described, 1-2–1-4

F

function descriptors
in the Attribution of CPU Execution Time report, 3-9, 4-4
in the Attribution of CPU Wait Time report, 4-2
in the Program Usage by Procedure report, 3-4
in the STROBE Performance Profile, 1-1, 1-5
in the Wait Time by Module report, 3-15

I

I/O Facility Utilization Summary report, 3-17
IDMS. See CA-IDMS
IDMS=NOADSO operand, 2-2
IDMSINTB (Batch Interface)
in the Attribution of CPU Wait Time report, 4-2
in the Wait Time by Module report, 4-2
wait time reported to, 4-1
IDMSSVC, STROBE-generated module name, 1-5
IDMSXXX, STROBE-generated module name, 1-5
indexed sets, 4-5
Indexers
executing, 2-3
map data sets, 2-3
required parameters, 2-3
indexing
and ADS/O dialogs, 2-2–2-5
concatenating listing data sets, 2-4
example of a report with, 3-6
with STROBE/ISPF, 2-4
with the STROX procedure, 2-4–2-5
instance number, appended to a dialog, 3-5
invoking, STROBE CA-IDMS Feature, 2-1–2-2

M

map data sets
concatenating, 2-3–2-4
creating, 2-4
measurement requests, submitting, 2-1–2-2
Measurement Session Data report
explained for batch regions, 4-1
explained for online regions, 3-1–3-2
measurement time, session duration, 1-1
measuring, sample data sets, 1-1
modules, STROBE-generated naming convention for, 1-5
Most Intensively Executed Procedures report
compared to the Attribution of CPU Execution Time report, 3-10
described, 3-2

N

naming conventions
for CA-IDMS nucleus and application programs, 1-1
for STROBE-generated module names, 1-5
for system service tasks, 1-1
pseudo-taskcodes, 1-5
NOADSO operand, 2-2
NOATTR, keyword parameter, 2-6
NOATTR=IDMS operand, 2-6
NOIDMS keyword parameter, 2-1

P

panels
TAILOR REPORTS, 2-5
Performance Profile
ADS/O, dialog reports, 3-5–3-8
ADS/O, dialog statement text included in, 3-8
CPU execution reports in, 3-2
described, 3-1
overview of I/O reports, 3-16
program module reports, 3-4–3-5
transaction reports, 3-11
wait time reports, 3-15, 4-2–4-4
program module reports, 3-3–3-4
Program Section Usage Summary report
explained, 3-3
Program Usage by Procedure report
compared to the Attribution of CPU Execution Time report, 3-10
explained, 3-4
for CA_IDMS-DC/CV mode, 1-3
pseudo-modules, .SYSTEM, 1-5, 3-3
pseudo-sections, .IDMS, 1-5
pseudo-taskcodes, described, 1-5

R

Reporter, 2-3
REPORTS=ALL, ADSORPTS parameter, 2-3
Resource Demand Distribution report
  explained, 3-16

S

sets, effects of organization, 4-3, 4-5
sorted sets, decreasing the search time, 4-4
statement offsets, 4-3
STROBE Advanced Session Management Feature, 3-14
STROBE CA-IDMS Feature
  executing the ADS/O Indexer program, 2-3
  I/O reports, 3-16
  invoking, 2-1–2-2
  measuring batch regions
    Attribution of CPU Execution Time report, 4-4
    Attribution of CPU Wait Time report, 4-2
    Measurement Session Data report, 4-1
    Wait Time by Module report, 4-2
  measuring online regions
    attribution reports, 3-9
    CPU execution reports, 3-2
    dialog reports, 3-5
    I/O utilization reports, 3-16
    program module reports, 3-3
    transaction reports, 3-11
    Wait Time by Module report, 3-15
  using, 1-1
STROBE-generated module name
  IDMSSVC, 1-5
  IDMSXXX, 1-5
STROBE/ISPF
  indexing, 2-4
  invoking CA-IDMS data collection, 2-2
  limiting attribution reports, 2-5
  suppressing
    all attribution reports, 2-6
    CA-IDMS attribution data collection, 2-1–2-2
    CA-IDMS attribution reports, 2-6
STROE procedure
  suppressing attribution reports with, 2-6
  tailoring attribution reports with, 2-5
STROX procedure
  required parameters, 2-4
STROXE procedure
  suppressing attribution reports with, 2-6
  tailoring attribution reports with, 2-5
STRXADS, executing, 2-3–2-4
SUBSYSTEM field
  in the Measurement Session Data report
    for batch regions, 4-1
    for online regions, 3-1
system modules
  identifying activity within, 1-5
  identifying heavy CPU consumption, 3-2
system service tasks
  CA-IDMS-DC services, 1-2
  identification of, 1-5
  naming convention, 1-1

T

TAILOR REPORTS panel
  discussed, 2-5–2-6
tasks
  CA-IDMS-CV tasks, 1-3
  CA-IDMS-DC tasks, 1-2
  identification of system tasks, 1-5
  naming system service tasks, 1-1
transaction reports, 3-11–3-15
Transaction Summary report
  described, 3-2
  explained, 3-11–3-12
Transaction Usage by Control Section report
  described, 3-2
  explained, 3-12–3-15
  identifying performance improvement opportunities with, 1-3
Transaction Usage Summary report. See Transaction Summary report

U

Universal Communications Facility (UCF), 1-4
utility programs, ADSORPTS parameters, 2-3

V

VIA, using as a set key, 4-5

W

wait time
  DEBUG DML option requirement, 2-2, 3-9
  reports
    for batch regions, 4-2–4-4
    for online regions, 3-15
Wait Time by Module report
  described for batch regions, 4-2
  described for online regions, 3-15
wait time, identifying
  sites of wait
    in batch regions, 4-2
    in online regions, 3-15