

1. Introduction

Cypress was motivated by the need for database management facilities in Cedar. Its goals are (1) convenience, making it easy to build new database applications in Cedar, and (2) reasonable performance for a wide class of applications. If you are unsure whether an application you have in mind is appropriate, talk to someone who has built one. You should also consult others to see what auxiliary documentation and folklore is available.

The conceptual data structure provided by a database system, analogous to the Mesa type system, is called a *data model*. The data model specifies how data may be structured and accessed; thus the data model is what you need to know in order to use the system. The Cypress data model described herein includes features that are accepted in some form in a number of models in the literature, particularly those that seemed most useful for database applications we envision in CSL.

The data model is described conceptually in Section 2. In Section 3 we describe the Cedar interface to Cypress: DB.mesa. Finally, Section 4 provides a simple example of Cypress use. We cover only the abstract model and the DB interface, not its implementation or rationale. The reader interested in these should consult CSL Report 83-4. A set of formal axioms for the model can also be found there.

Clients interested in tools for examining or modifying databases, either as an end-user or for debugging an application program, should consult the database tools documentation:
[Indigo]<Squirrel>Documentation> SquirrelDoc.tioga, .press.