Time Series Database Interface*

Paul D. Gilbert
Department of Monetary and Financial Analysis, Bank of Canada,
234 Wellington Street, Ottawa, Canada, K1A 0G9
pgilbert@bank-banque-canada.ca

Abstract

This presentation describes a package that abstracts an interface to time series databases, and a related group of packages that implement interfaces to SQL databases and to Fame through the PADI protocol. The TSdbi package, which implements the abstraction, imports the DBI namespace and DBI functions that support many SQL databases. For these cases there is limited need for code specialized to the specific database. This has been implemented in packages TSMySQL and TSSQLite, which require packages RMySQL and RSQLite respectively. It should also be possible to use the abstraction with RODBC (in progress but untested at the moment). TSdbi can also be used to interface to other time series databases, but this will typically require more database specific code below the abstraction. A working interface to Fame is implemented in the package TSpadi. It should also be possible to implement a more direct connection using the fame package.

Time series databases are typically simple in the sense that series are named with a unique identifier, and queries are limited to lookups using this key. From this perspective an SQL database is hardly needed. Apart from the abstraction, which is useful to make other code independent of the database implementation, the the main advantages are to use the database's client/server protocol, the ability to handle endian issues, and security features. However, when an SQL database is used, additional features can be added: it is possible to have vintage and panel dataset, with the same identifier used for different release dates and/or different panel members.

The package also (potentially) allows choice of the R representation to use for the time of the series data points. The default is ts where applicable, and zoo otherwise. (At the moment, only the default is working.)

The structure of the back-end SQL data bases and some utilities for implementing them will also be discussed.

*The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.