Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Engineering & Technical Services Division

CUSTODY/OWNERSHIP DICHOTOMIES IN THE U.S. PETROLEUM SUPPLY SYSTEM

David F. Cahn

October 1980

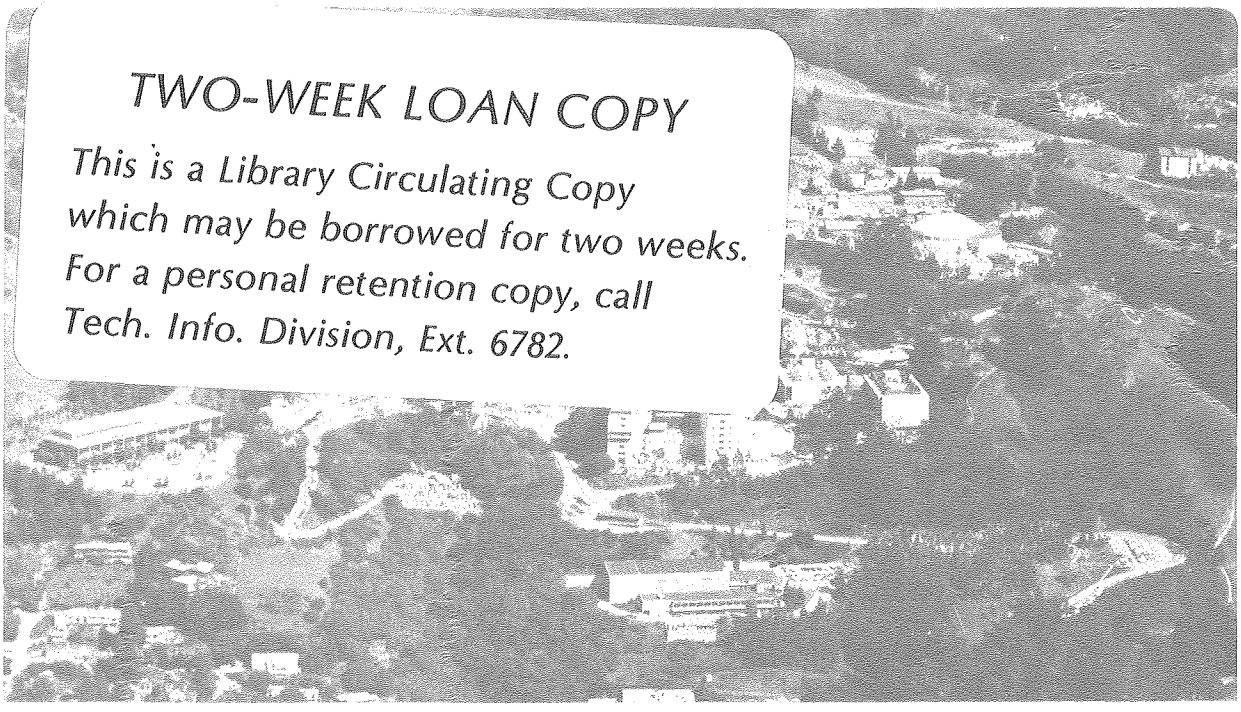

Prepared for the U.S. Department of Energy under Contract W-7405-ENG-48

DISCLAIMER

LBL-10583

CUSTODY/OWNERSHIP DICHOTOMIES IN THE U.S. PETROLEUM SUPPLY SYSTEM

David F. Cahn
Information Methodology Research Project
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

October 3, 1980

Inherent operating characteristics of the U.S. petroleum industry make it difficult to track petroleum ownership and custodial holdings simultaneously, and yet this is necessary for performance of various mandated DOE functions. In petroleum emergency management, for example, a common goal--compensation for regional shortages of specific products--requires information on existing inventories and available transportation mechanisms for moving them to the affected region. This information is custodial in that it concerns total inventories at a given site regardless of ownership. However, if federal action, such as a refinery yield order, is taken, this action must be directed at an existing legal entity (such as a corporation), and that entity can only be expected to comply as regards its own products. In other words, if the company is refining a run of someone else's crude oil (frequently the case in the industry), it would feel understandably reluctant to vary the product mix without a direct order from the legal owner of the oil. Thus, to be effective, a refinery yield order must propagate down an ownership chain to affect oil custodially held at a given refinery.

Petroleum Industry Practices

The manner in which business is conducted in the petroleum industry determines the traceability of both ownership and custodial information. A prime factor in the industry is that the relative difficulty and expense of moving crude oil and petroleum products has led to a system of 'swapping' among the participants, and, consequently, to the divorce of petroleum transactions from petroleum movements. Let us say, for example, that Company A has 100,000 bbl. heavy crude surplus in one of its Los Angeles refineries, but needs 50,000 bbl. of heating oil in Chicago and 10,000 bbl. of gasoline in New York to supply local demand beyond its present production in these areas. Rather than move the surplus crude to its own refineries in Chicago and New York, Company A looks around for a Company B that has need of an additional 100,000 bbl. of crude locally in Los Angeles and can supply Company A with its additional refined product needs in Chicago and New York. A contract is negotiated between A and B, and everyone's needs are optimally met without need of a major transportation expenditure.

In fact, several available latitudes can complicate the situation. First of all, the contract is very often less than formal: it is often advantageous to swap oil rather than dollars (for tax and other reasons), and so imbalances are made up with i.o.u.'s. Secondly, the transaction often involves more than two parties (``I'll give you x bbl. that I have at Peoria and y bbl. that Company C owes me and can supply at Evanston.''), and utilizes i.o.u.'s that have been held for some time. Third, since the trades cross product lines, a given refinery may instantaneously be holding crude oil that is to be refined into gasoline that has already been commit-

ted to some other company; alternatively, the same refinery may be holding crude oil that it is merely under contract to refine and in which it has no equity. Of course, it might subsequently acquire equity, via a trade, in the heating oil it produces from this crude oil. Fourth, while among independent companies it is reasonably certain that some form of formal accounting procedures trace the swapping activity [1], it is not at all clear what external evidence, if any, would be visible for trades entirely within any of the large, vertically-integrated oil companies, or among subsidiaries of a single conglomerate.

The complications introduced by these kinds of multiple brokerage make it extremely difficult for any outside agency to trace oil ownership, or to link ownership to custody. Clearly, oil in storage at one plant can change hands a dozen times wihout any physical movement, and it may be entirely unrealistic to expect the custodian to know who it belongs to at any instant. Furthermore, the argument can certainly be made (although there is sufficient impetus for these practices without it) that this form of multiple brokerage is also an excellent mechanism for hamstringing the federal price-tiering regulatory structure: obviously, after two or three swaps, it becomes difficult indeed to determine what wells a given parcel of oil came from, or, indeed, whether it is even domestic.

Even when oil does move physically, the picture is not clear. There is, for example, no universally accepted definition for the instant of equity change, and physical transfer and accounting transfer may differ by a month or more. For example, Company A reports a parcel sold on the date shipped but Company B reports it bought on the day it <u>pays</u> for it. The parcel may actually reach Company B in two days, but it may not be paid for

until a month later. In a monthly reporting scheme, there may be no provision for linking parcels that are reported in different sampling intervals (thus vastly complicating validation), and it is further apparent that the parcel can literally disappear from the system for a month or more at a time. If, instead of a single transaction, A and B engage in a multiple transaction (perhaps with an intermediary Company C brokering the oil), the linkup problem becomes overpowering to any attempts at tracing either ownership or flow.

A further list of known anomalies to fill out our Pandora's box must include the following. First, it is frequently unrealistic to expect even the participants to have more than local information; in other words, an individual actor can only know from whom he purchased a parcel and to whom he sold it, and any data requests concerning the more distant parts of the transaction chain will necessarily elicit speculative answers. Second, oil loses its parcel identity when it flows into a tank, so a 'parcel' is really no more than a paper concept useful in transacting business; furthermore, higher aggregations (such as monthly totals shipped or sold from A to B) are artifacts of the reporting procedures and have very little connection to the actual manner in which business is conducted. These two points help to explain the constant revision to which petroleum data appear subject: 'ultimate consignee' (whether specified as company, SMSA, state, or PAD) is at best a highly speculative thing to ask a company to report, and the cumulation of data on a monthly basis forces generalizations on the reporting companies that are unrelated to their normal business data.

Third on this 'miscellaneous' list, the transportation entities (especially the pipeline companies) serve several roles in the system. They

sometimes act as contract agents, in which case they are custodians, but not owners, of the oil they are carrying; in other instances, they may buy oil from a universe of suppliers and market it to a separate universe of customers, thus serving as brokers themselves. Lastly, jobbers and many of the brokerage entities may own oil on paper only, bartering vast quantities of oil without having any physical facilities whatever. Such holding companies are spectral at best, and greatly complicate an already intricate business network.

Query Classes

A number of specific aims characterize expected federal usage of petroleum data. In turn, these have specific data dependencies.

A most important distinction is the one noted in the introductory paragraph, namely that between monitoring and intervention. For a large class of emergency management and policy decisions, custodial information is most appropriate. A physical stock view provides the clearest image of shortfalls and impact propagations, and is useful for recognizing both long term fuel dependencies and their modification. For other types of monitoring tasks, however, concerning such matters as allocation fairness, fair business practices, and numerous regulatory matters, ownership information is essential.

For any intervention activities, there is little choice: it is necessary (1) to be able to hold some legally-defined entity accountable for compliance with the federal directives, and (2) to direct federal orders to those entities that have jurisdiction over the substance of interest. The first consideration dictates an approach down the corporate chains to the

site at which action is possible, and the second requires that the chains followed be those of the product owners rather than the custodians.

Common Denominators and Data Perspectives

A number of similarities and differences characterize the custody and ownership viewpoints. A very useful perspective, as diagrammed in Figure 1, is that custody and ownership are merely two of a number of aggregation projections possible on the same pool of raw data. An essential conclusion of this viewpoint is that one must be extremely careful in defining the basic level of raw data, since it is only at this rudimentary level that the two views can be reconciled.

To illustrate this point, consider the 'shipment' as a basic unit. One can aggregate shipments to form weekly or monthly totals flowing among geographic or corporate entities. As long as we always tally comparable data at any level, our physical flow aggregations should yield a consistent picture of actual stocks and movements. We may choose to tally by corporation on either an ownership or a custodial basis, but we must be careful not to mix the two if we wish to avoid doublecounting and miscounting errors.

Similar to physical shipment tallying, we can tally sales transactions by corporation, and can aggregate to any level. However, again, we must be careful always to obtain only data on sales, not on physical delivery.

At the level of 'shipments' or 'transactions', the data concern different types of activity that can befall a parcel of oil, and, because the activities differ, they must be kept isolated. If, however, we adopt an

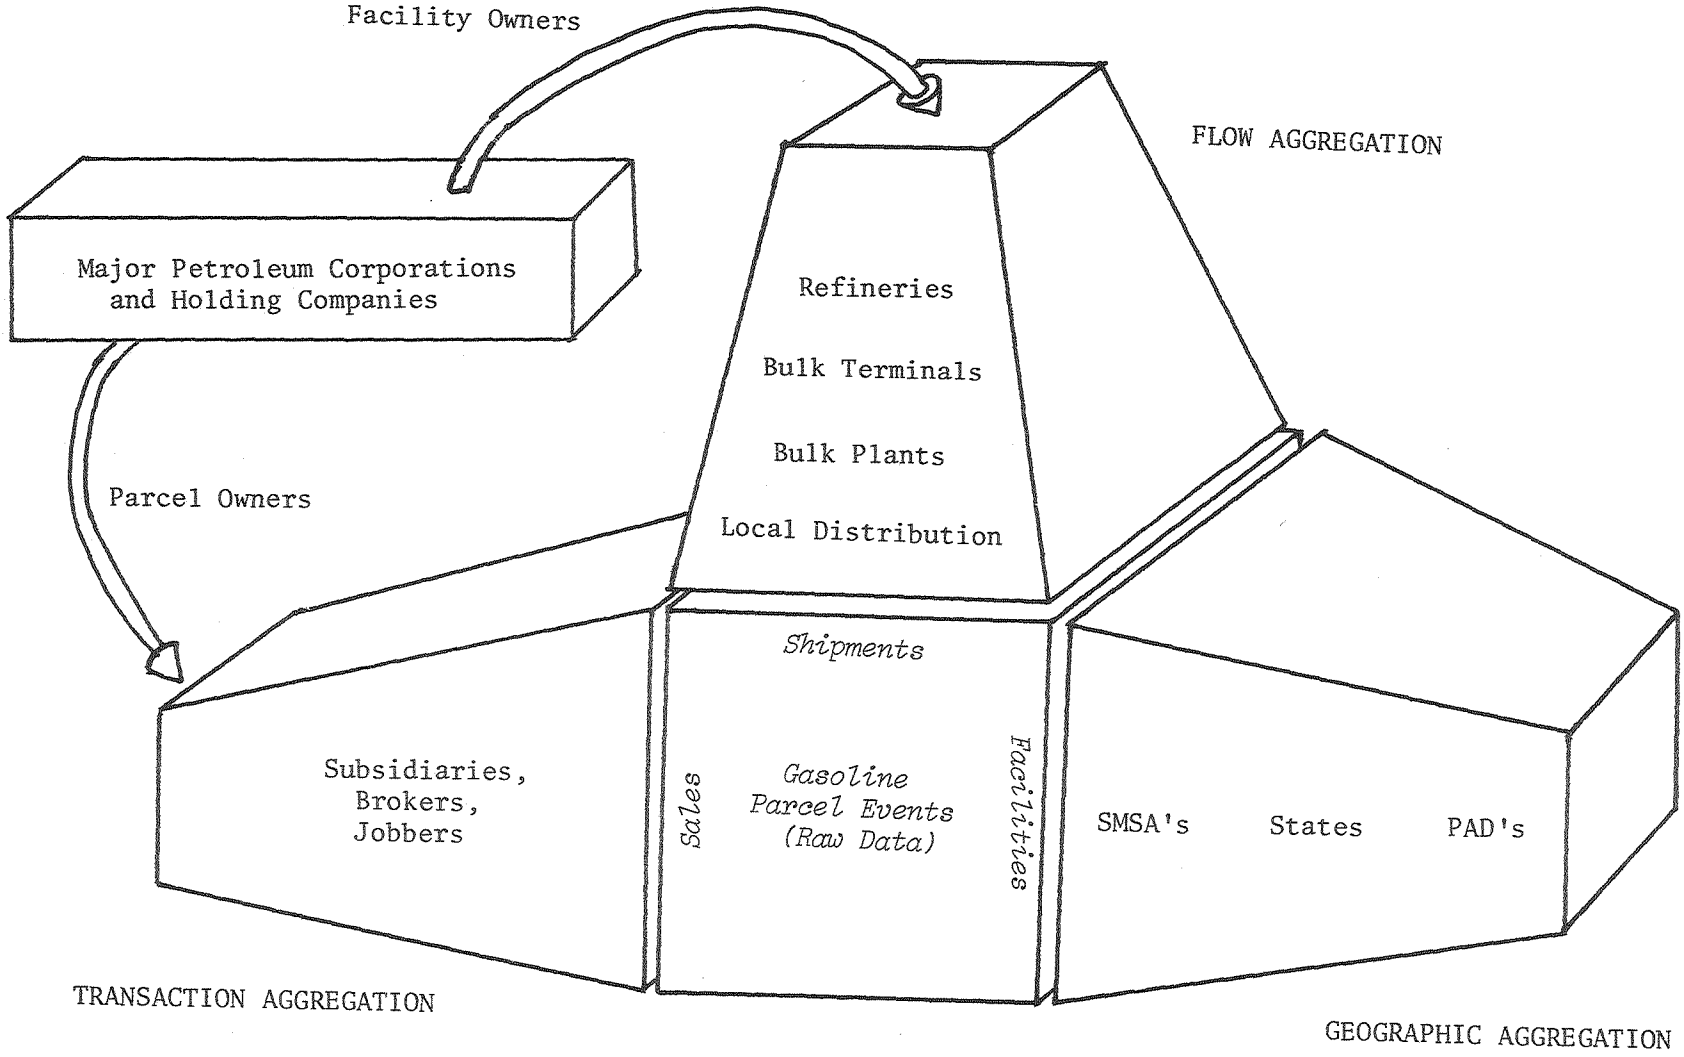

Figure 1. Custody, ownership, and geographic aggregation projections based on an 'event' infrastructure.

even more basic view, and deal with generic 'events' that can befall an oil parcel, then we can define a state vector that mirrors both ownership changes and physical changes. If there is an 'owner' variable in the state vector, a change in ownership is merely an incremental event befalling the parcel; the change modifies the 'owner' variable only, and all other state variables remain constant. The same is true in the physical domain if we declare a 'location' variable that identifies, say, a particular refinery. The refinery entity carries corporate ownership as well as physical location information, so we have a custodial tie as well as a geographic one. At the event level, a sale changes the ownership variable without changing the physical location; a shipment has the inverse effect. The 'event' viewpoint, however, allows both views to apply to a 'event' viewpoint, however, allows both views to apply to a single informational entity, a tie that fails at any higher level for lack of convergence. The implication, namely that data collections based on higher aggregate levels can reconcile neither custody and ownership nor physical flow and transactions, explains much of the confusion that has come to surround these dichotomies. An appropriate data perspective (in this case, level dependent) is crucial to the ability to approach the problem.

A second view of custody and ownership that has potential utility is diagrammed in Figure 2, using refineries as an example. It is possible to explicitly relate sectors of the data pool in such a way that the orthogonal projection aspects of custody and ownership become clear. In the matrix of Figure 2, the columns contain data on each company's custodial holdings and the rows contain data on each company's equity. An individual data point naturally participates in both row and column summations, and thus the real situation is mirrored appropriately. Sectors of the matrix

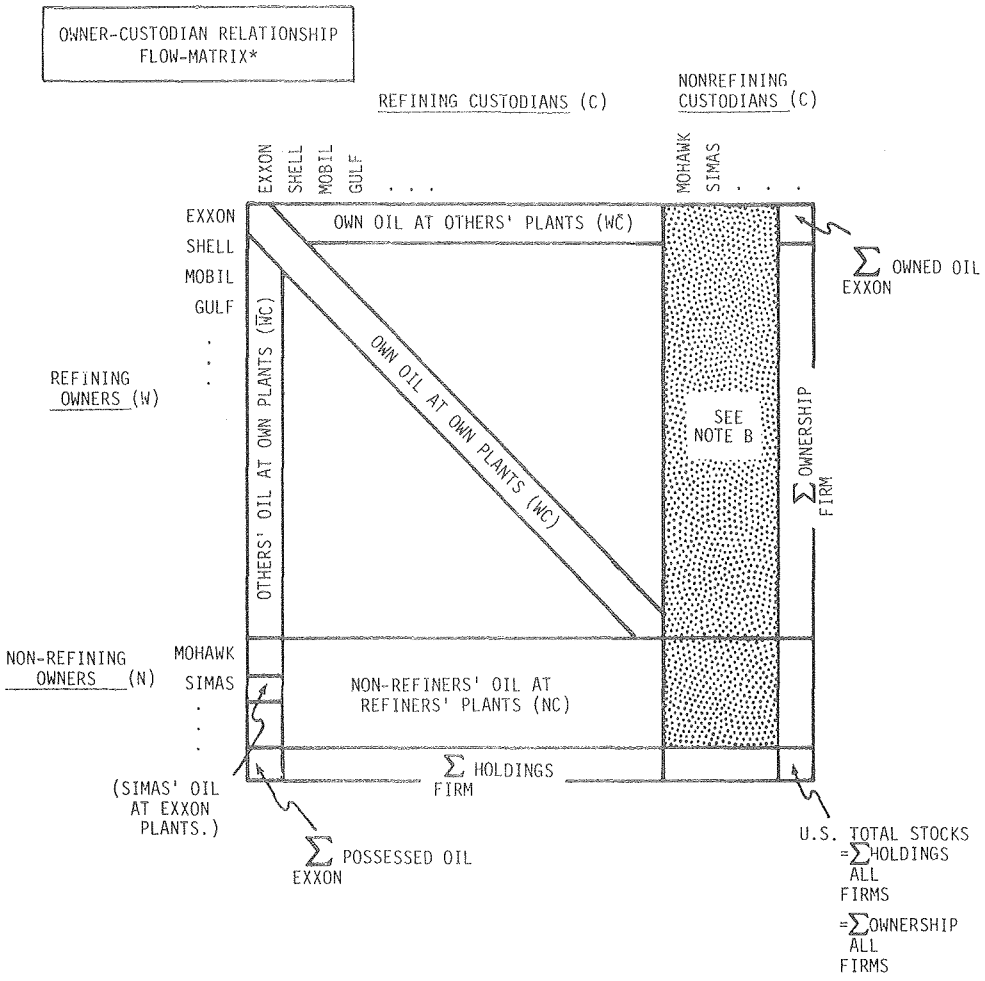

*NOTE A: MATRIX ENTRIES CAN BE FIRM TOTALS (AS PRESENTLY REPORTED), OR CAN BE FURTHER SUBDIVIDED BY INDIVIDUAL PLANT, STATE, ZIP, ETC., AS DESIRED:

EXXON [PLANT 1, PLANT 2, PLANT 3]

or EXXON [CALIF, TEXAS, NJ] etc.

**NOTE B: UNDER DEFINITIONS OF PRESENT FORMS, NON-REFINERS DO NOT HAVE THEIR OWN PLANTS. HOWEVER, IF 'PLANT' IS DEFINED TO INCLUDE ANY STORAGE OR OTHER HOLDING FACILITY AS WELL AS REFINERIES, THEN THIS PORTION OF THE MATRIX WOULD CONTAIN INFORMATION ON OIL HELD FOR RESALE BY NON-REFINING BROKERS (OFF-DIAGONAL PORTION) AS WELL AS OIL HELD AND OWNED BY NON-REFINERS (DIAGONAL PORTION).

Figure 2. Custody/ownership relationships among refineries, sectored to correspond with EIA reporting categories.

reflect the range of possibilities:

- the major diagonal contains all entries where the custodian and owner coincide (one data point per company);

- the row for each company references all its oil equity around the country, and this row may be summed to indicate its total ownership;

- the column for each company references all its custodial holdings both of its own oil (the single diagonal entry within the column) and of others' oil, and may likewise be summed to provide a custodial total for the company;

- non-refiners (i.e., companies that do not operate their own refineries and contract for refinery service at other companies' plants) account for a sizeable fraction of the oil in the system, but this oil is held custodially at others' facilities, hence the distinct matrix sector allocated to them.

A matrix such as that shown is reasonably robust over aggregation level of the contained data points, although it is still governed ultimately by the reconciliation arguments previously discussed; the data points can refer to individual plants or to corporate summations with equal validity, although geographical traceability, is an obvious sacrifice as the level is raised. Similarly, the matrix entries can be collected as time cumulations or snapshots with equal validity, although simultaneity of their collection must always be preserved. The limitations of this view, however, relative to the reconciliation arguments raised previously, must be kept in mind: a view such as this is adequate to capture overall custodial holdings and equity, but is not adequate to trace either buy-sell chains or consignment-delivery chains to their root levels. Nonetheless, Figure 2 does lend some organizational perspective to the field by indicating some straightforward ways in which the custodial and equity views are related.

Approaches and Methodology

Given the arguments that have been made, one can state that many of the reconciliation difficulties that seem to plague the custody/ownership problem are artifacts caused by inappropriate data views. An appropriate state space, predicated in this case on 'events' as a base grain, allows common treatment of both domains and reveals the phenomena of interest in tracing the various functions of the petroleum industry. Higher aggregation levels then become derivative on the 'event' data; downward speculation is eliminated at the same time that validation checks become possible. Otherwise untraceable loopholes (such as multipele brokerage) can no longer occlude the system and destroy its tracing effectiveness. The penalty, of course, is in the number of data entries that must be processed, but arguments made elsewhere [2] suggest that the extra burden may be artificial.

The base set of state variables must span the range of possible events if the state space view is to function effectively. In other words, the prime requirement is that no event go undetected; any status change that can affect a parcel of oil (ownership, location, physical composition, etc.) must be uniquely represented. To the extent that the state variables selected are independent, the number required to span the event spectrum will be minimal. (in earlier work in this area [3], we found a set of twelve state variables that met the sufficiency criteria at the event level.)

In designing data collection mechanisms for both custody and ownership domains, it is important to ask each reporting entity only for data it can report with certainty. In most cases, this means that the operator of a given facility should report no more than his immediate source of the parcel and the entity to which he transferred it. In the physical flow

domain, this represents the immediate supplier and consignee; in the transaction domain, it represents the immediate seller and customer. Since more distant state changes in the event chain are not within his direct purview, he should not be asked to infer them. Rather, the system should be configured to treat all elements (refineries, pipeline companies, brokerage companies, etc.) equivalently, so that the event chains can be coupled after the fact. Particularly, a refiner in Louisiana should not be asked how much gasoline he sold for ultimate use in New England; instead, he should only report how much he sold to the owners of the Colonial Pipeline, and they should report the volume they sold to New England. In this regard, it should be noted that a physical flow network is useful in defining the range of possibilities for transaction chains, but is not conclusive. The transactions, by virtue of their independence of physical movements, may follow entirely different lines than those suggested by flow patterns.

In tackling the custody/ownership problem, many of the same design qualities that aid flow visualization have importance:

- Mechanistic uniformity of the data structure

- Consistency and simultaneity of the data collected

- Comprehensiveness of the data collection at the aggregation level selected

- Isolability of expected systemic errors; ability to restrict their propagation

- Support of validation methods. (Generally, three levels of validation exist: (a) multiple entry, which allows direct comparison among reports from different respondents; (b) internal validation based on locally limited knowledge of the data pool; (c) external validation, or crosschecks based on general knowledge of the system's operating characteristics.)

In all cases, it should be possible to define ideals and compromises, and

to identify levels of acceptable performance. The key generally lies in approaching the problem from a direction in which its internal workings are adequately exposed, and, in the case of the custody/ownership dichotomy, this optimal approach direction is quite selective.

References

[1] Peterson, J.L., et al., "Petroleum Products Systems Evaluation Task Force Revised Report." DOE/EIA, Office of Energy Operations, December 1979.

[2] Cahn, D.F., "Common Petroleum Emergency Characteristics and EEMIS Aids for their Management." DOE/EIA Office of the Energy Emergency Management Information System, February, 1980.

[3] Krishnan, V.V., and Cahn, D,F., "Material-Flow Data Structures as a Basis for Energy Information System Design." LBL-10248, to be presented at ASIS Mid-Year Conference, Pittsburgh, May 1980.

Acknowledgements

This work was supported by the Engineering and Technical Services Division, Information and Data Analysis Department of the U.S. Department of Energy under Contract W-7405-ENG-48.