SANDIA REPORT
SAND2012-9487
Unlimited Release
Printed November 2012

A Demand-driven, Capacity-constrained, Adaptive Algorithm for Computing Steady-state and Transient Flows in a Petroleum Transportation Network

Walter E. Beyeler
Thomas F. Corbet Jr.
Jacob A. Hobbs

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2012-9487
Unlimited Release
Printed November 2012

A Demand-driven, Capacity-constrained, Adaptive Algorithm for Computing Steady-state and Transient Flows in a Petroleum Transportation Network

Walter E. Beyeler
Thomas F. Corbet Jr.
Policy and Decision Analytics
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-1138

Jacob A. Hobbs
Computer Science Department
University of New Mexico
Albuquerque, New Mexico 87106

Abstract

We developed an algorithm to perform simulations of a supply network for crude oil and refined products in order to estimate the consequences of disruptions to components of the network. Components include oil fields, import terminals, refineries, transmission pipelines, tank farms, and distribution terminals. The physical system is represented as network connections, capacities, and inventories.

The governing equations describe mass balance in a non-linear diffusive system in which flows in the network are along gradients in a potential field. Each node in the network has a defined storage capacity and desired storage amount. The potential at each node is a function of the difference between the actual and desired amount of fluid stored. The potential can be thought of as the balance between the desire to increase inflows to maintain the desired storage level and the willingness to provide fluid for consumption or outflow to downstream nodes.

INTRODUCTION

Our objective in developing this algorithm is to perform simulations of a supply network for crude oil and refined products in order to estimate the consequences of disruptions to components of the network. Individual components of this network represented by the algorithm include oil fields, petroleum (crude oil and refined products) import terminals, refineries, transmission pipelines, tank farms, and distribution terminals.

Three high-level requirements for the algorithm are:

Flows of petroleum over the network are driven by the demand to consume refined products aggregated at distribution terminals.

> The algorithm strives to provide the amount of fuel normally consumed at distribution terminals. Shortages of fuel occur if fuel is physically limited.

Flows of petroleum are constrained by the capacities of network components.

> Network components that have capacities include crude oil production from oilfields, transmission pipelines, refineries, tank farms, and terminals for water-borne shipments of petroleum. Reduced capacity of any of these components could contribute to fuel shortages at distribution terminals.

Flows of petroleum adapt to minimize the impacts of disruptions to network components.

> The simulated response to a disruption to any component includes dynamic rerouting of shipments, drawdowns of inventory, and using surge capacity in refining, pipeline shipments, and petroleum imports to mitigate the impacts of a disruption.

The ability of an actual petroleum network to respond to disruptions depends on both the capacity of the engineered system and human business and operational decisions. Selecting an approach to best simulate this mixed human- physical system for the intended use is challenging.

Given that engineering models of petroleum infrastructures that solve fluid dynamic equations are readily available, it is tempting to use these models to represent the physical part of the mixed system. Our experience is that this approach is problematic for three reasons. First, it is not practical to build models of large networks (for example, the entire U.S.) because the engineering details of system components are not available outside of the private companies that own and operate them. Second, if such a model could be constructed it would be overly complicated and expensive. The third reason is that an engineering model of a petroleum network would have to be “operated” by an analyst or an algorithm. For example, if a business decision is made to increase flow in a pipeline, an engineering decision must be made to adjust a pump rate to accomplish the flow increase. The instruction to adjust the pump rate would have to be provided to the simulation. For disruption analyses, we find more value in simulating business decisions than engineering operations because the business decisions drive the operations. Consequently, this algorithm represents the physical system as network connections, capacities, and inventories rather than engineering descriptions of components. Flows on the

network are driven by simple business rules that can be adjusted to represent different assumptions about operations or policies.

The basic assumption of the business rules is that the system, through market behavior and common business practices, will use all available capacity to minimize fuel shortages to customers. Therefore simulated fuel shortages (defined as consumption less than during undisturbed conditions) indicate that the simulated network is constrained in some way such that it cannot provide all of the fuel normally consumed at that location. In this sense the simulation predicts the best plausible outcome.

The governing equations describe mass balance in a non-linear diffusive system in which flows in the network are along gradients in a potential field. Each node in the network has a defined storage capacity and desired storage amount. The potential at each node is a function of the difference between the actual and desired amount of petroleum stored. The potential can be thought of as the balance between the desire to increase inflows to maintain the desired storage level and the willingness to provide fluid for consumption or outflow to downstream nodes.

Flows on links are a nonlinear function of the potential difference across the link and link capacity. The function assures that flow approaches zero at very small gradients of potential and approaches the maximum capacity at large gradients. This function assures that flows on the network are constrained by capacity.

Consumption at each distribution terminal is a function of the potential at that node. That is, consumption is a potential-dependent flow on the network boundary. It is a proxy for the effect of price on demand. Consumption declines as storage at the distribution terminals falls below its desired level. Inflows to the network are represented as fixed potential boundary conditions connected to the network by a link with fixed maximum capacity. This combination of boundary conditions assures that flow in the network will be driven by demand for consumption.

Finally, the consequence of a disruption should include the time it takes for the system to replenish inventories of petroleum after the disruption ends. Defining the potential at each node as a function of the difference between the actual and desired amount of petroleum stored assures that the simulated system will return to its pre-disruption condition at some time after the disrupted component is repaired.

NETWORK FLOW EQUATIONS

Network flow is described by a diffusion equation with non-linearities introduced to model capacity limits in the system. The system state is defined by a potential function $\{s_i(t)\}$ defined at the network nodes i. A link connecting nodes *i* and *j* allows flow between the nodes from the node with higher potential to the node with lower potential. Flow rates are limited by a (directed) capacity associated with each link, c_{ij}. Assuming $s_i > s_j$, the flow from node *i* to node *j* is given by:

(1)

where u_{ij} is a utilization parameter and the function f(x) models linear resistance as x->0 and enforces the capacity limit for large x:

(2)

Figure 1. Graph of Equation 2.

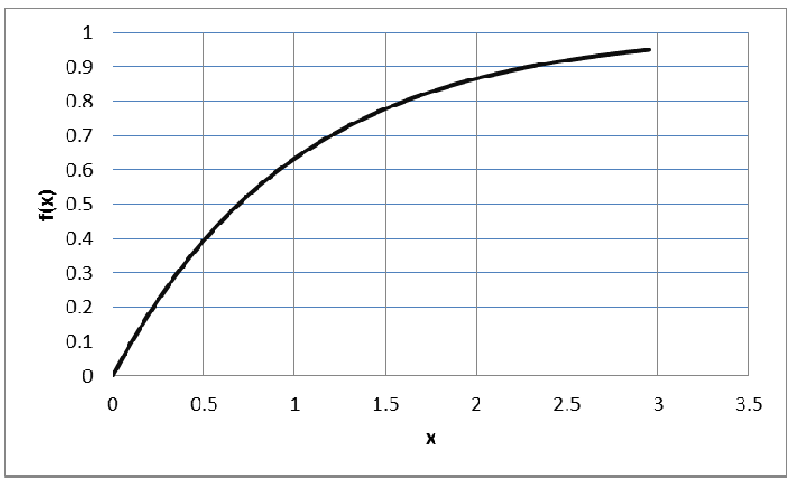

Flow through the network is driven by a set of specified withdrawal rates and a set of source nodes with constraints on their inflow rate. Source behavior at a node *k* is modeled as a fixed potential connected to node k through a capacity-limited link. Source flow into node *k* is then simply given by equation (1):

(3)

The fixed potential may be the same for all sources in the model, or may be specified separately for each source. In either case, the actual flow rate from the source into node *k* can be adjusted by changing the utilization factor .

Demand at a sink node l is characterized by a desired flow rate d_l. In general we assume that the baseline system capacities can satisfy the set of demands, so that there is a feasible solution to the set of equations for node potentials. In equilibrium, the net flow at each node i is 0:

$$\sum_j q_{ji} + q_{si} - d_l = 0 \quad \forall i \tag{4}$$

The equilibrium solution $\{\hat{s}_i\}$ is obtained by solving equations (1-4).

When the network is disrupted, the baseline demand flows d_l may not be achievable. The system should allow for demand flows to be reduced as stress in the system increases. We do this by replacing the fixed-flow demands with fixed potentials in the form of (1) whose parameters are derived to satisfy the baseline demands. Each demand node l has an associated relative surge consumption $\epsilon_l > 0$ and a demand utilization u_{dl}. The capacity of the new demand link is then $d_l(1+\epsilon_l)$, . Because the flow rate is (by definition) d_l, the fixed potential for the demand node, s_{ld}, must satisfy:

$$d_l = d_l(1+\epsilon_l)\, f((\hat{s}_l - s_{dl})u_{dl}) \tag{5}$$

Using (2) for the relationship between potential difference and relative flow,

$$s_{dl} = \hat{s}_l - \frac{1}{u_{dl}}\left[\frac{1}{(1+\epsilon_l)^2-1}\right]^{1/2} \tag{6}$$

Each node has some capacity to store fluid, and the stored volume v_j adjusts to transient imbalances in flow induced by changes in the system structure or in external stress. Transient equations for $\{s_i(t)\}$ reflect the following assumptions:

- The equilibrium state $\hat{s}_i$ is associated with a stored volume equal to some specified fraction $\beta_i$of the node's storage capacity v_i^T. This fraction can vary across nodes but will generally be specified as a global constant.
- Small variations in $s_i(t)$ reflect changes in stored volume proportional to the node's capacity.
- Stored volume at a node can never be less than 0 and never be greater than the node's capacity.

The first two assumptions suggest a linear relationship between s_i and v_j near equilibrium, while the last shows that nonlinearities are needed to deal with limiting conditions. The properties of the required function are sketched below:

Figure 2. Dependence of stored node volume on state variable S_i.

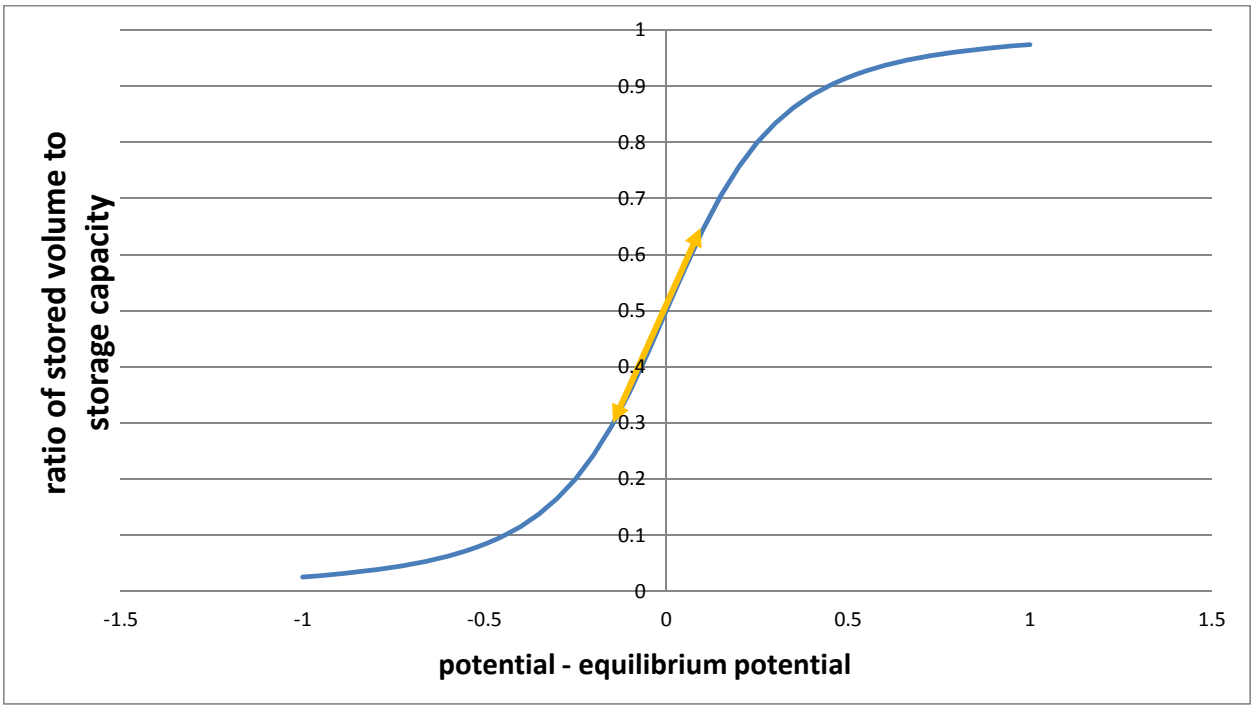

A simple function honoring the desired constraints is:

$$\frac{v_i(t)}{v_i^T} = \frac{1}{2}\{1 + \frac{\frac{s-a}{b}}{\left[1+(\frac{s-a}{b})^2\right]^{1/2}}\} \tag{7}$$

Where a and b can be expressed in terms of the target fraction β_i and the storativity γ_i

$$b = \frac{1}{2\gamma}[1 - (2\beta - 1)^2]^{1/2} \tag{8}$$

$$a = \hat{s}_i - b\left[\frac{(2\beta-1)^2}{1-(2\beta-1)^2}\right]^{1/2} \tag{9}$$

Parameters β_i and γ_i do not figure in the equilibrium solution, and may be assigned independently from the basic network definition. This separation permits alternative policies for baseline inventory and willingness to release inventory to be evaluated without having to repeat derivation of the baseline potentials.

In the transient case, net inflow into a node results in the accumulation of stored fluid, and the dynamic equations for the state variables can be derived from mass conservation:

$$\sum_j q_{ji} + q_{si} - d_l = \frac{dv_i}{dt} = \frac{dv_i}{ds_i}\frac{ds_i}{dt} = \frac{v_i^T}{2b}\left[1 + (\frac{s-a}{b})^2\right]^{-3/2}\frac{ds_i}{dt} \quad \forall i \tag{10}$$

Equation 10 along with (1-3) define the transient equations for s.

INPUTS AND OUTPUTS

The algorithm requires the network properties defined below. Some are used in the equilibrium solution, others in the transient solution. Some properties may have the same value for all network elements and so be specified globally.

Table 1. Node Properties

Name	*Description*	*Transient Only*	*Other Constraints*	*Optional Global Default*
s_{0k}	Source potential		Source nodes	Yes
c_{sk}	Source capacity		Source nodes	
u_{sk}	Source utilization		Source nodes	Yes
d_l	Sink (demand) rate		Sink nodes	
ϵ_l	Sink relative surge consumption		Sink nodes	Yes
u_{dl}	Sink utilization		Sink nodes	Yes
v_i^T	Storage capacity	Yes		
β_i	Storage target fraction	Yes		Yes
γ_i	Storativity	Yes		Yes

Table 2. Link properties

Name	*Description*	*Optional Global Default*
c_{ij}	Link capacity	
u_{ij}	Link utilization	Yes

Transient behavior is induced by changing the properties of one or more network elements.

NUMERICAL IMPLEMENTATION

The equations for the steady-state and transient flows can be readily formulated for solution by one of the standard numerical methods libraries such as the Apache commons math library (http://commons.apache.org/math/). However the nonlinearities in the resistance term (Equation 2) and storage term (Equation 7) complicate the solution of the flow equations. As link flows approach their capacities, for example, flow rates become decreasingly sensitive to small changes in potential. At small flow rates, $\frac{df}{dx} \rightarrow 1$, however where $f(x) = 0.95$, $\frac{df}{dx} = 0.05$. Such large contrasts in coefficients tend to make the flow equation matrix poorly conditioned, making convergence slow and subject to large local balance errors. There are several strategies for improving solution performance, such as local linearization of the flow equations around the current estimate for the solution, accelerating convergence by using local mass balance errors and approximations of the governing equations to find locally-valid analytical solutions, and by inflating link capacities to obtain an initial solution, which is iteratively used to seed a solution with capacities that approach the actual values. Direction-dependence of link capacities can potentially lead to instabilities in solution of the steady-state equations, especially if the capacities are updated between iterations on the basis of changes in apparent flow directions. We do not have enough experience to judge how often this might occur in practice, however establishing flow directions at the beginning of a solution cycle, and checking for consistency once convergence has been achieved, seems a reasonable precaution.

DISTRIBUTION

1	MS0899	Technical Library	9536 (electronic copy)

Sandia National Laboratories