Cover Page

Abstract ID	**8100**
Paper Title	Market-Based Linkages in Integrated Land Use Transport Models
Contact Author	John E. Abraham (primary author)
Email	jabraham@ucalgary.ca
Phone	1-403-230-5897
Fax	

MARKET-BASED LINKAGES IN INTEGRATED LAND USE TRANSPORT MODELS

John Edward ABRAHAM
Post-Doctoral Fellow
Department of Civil Engineering
University of Calgary
2500 University Drive NW
Calgary, Alberta
CANADA T2N 1N4
Tel: +1-403-230-5897
Fax: +1-403-282-7026
E-mail: jabraham@ucalgary.ca

John Douglas HUNT
Professor
Department of Civil Engineering
University of Calgary
2500 University Drive NW
Calgary, Alberta
CANADA T2N 1N4
Tel: +1-403-220-8793
Fax: +1-403-282-7026
E-mail: jdhunt@ucalgary.ca

Abstract: An economic market approach to urban system modelling is described, where all interactions are characterized as exchanges in a market. This leads to a natural partition of an integrated urban model into submodels based on the category of good or service being supplied or demanded, the type of agent making the demand or supply, and the time and place of interaction. Actors communicate through 6 defined operations on markets, decoupling the algorithms representing different behaviour. Agent based approaches to simulation are a natural extension of the market approach, allowing a transaction based simulation of heterogeneous spatial markets with individual agents making specific offers in specific places at specific times, and other agents accepting those offers at future specific times. Incorporating many existing modelling methods into such a framework requires a set of average prices by segmenting markets by commodity category, space and time, and wrapping the modules in interfaces that recast the inputs and ouputs into market operations. The paradigm is applied to two existing modelling frameworks: the Sacramento MEPLAN model and the Oregon statewide TLUMIP model.

Keywords: Land Use Transport Interaction Modelling; Model Theoretical Framework; Market Treatments

1. INTRODUCTION

1.1 Overview and Objectives

Modelling of spatial economic systems is a growing multidisciplinary. Urban economists, geographers, regional scientists, regional planners and transportation planners all have their own approach to the problem. Urban economics tends to treat the characteristics of space as a smooth mathematical function, and focuses on the properties of the mathematical function. Regional scientists have typically focused on larger regions, and tend to abstract space into a small number of "cities" or "regions". Regional planners focus on the land development process and its regulation. Transportation planners focus on the demand for transportation infrastructure and the transportation network's performance under congested conditions.

Within transportation planning, the classic modelling approach was originally non-economic, in that the demand for transport was characterized by simple regressions or physical analogy. More recently, in the past three decades, transportation modelling has adopted econometrics and statistics, to build more accurate and more defensible models about how individuals behave. Disaggregate modelling has proved to be especially powerful, since the diversity of individual characteristics is largely responsible for the diversity of travel demand. But transportation modelling is, for the most part, still very much transportation *demand* modelling. The only element of supply that is routinely considered is the fixed supply of transportation infrastructure and its performance under congestion.

Recently, however, *land-use transport interaction* models have become more important. These models focus on supply *and* demand, and hence markets, in final and intermediate goods and services, floorspace, labour, retail goods, and land. Since cities exist primarily to ease interaction, and a market is defined as an "institutional arrangement facilitating exchanges" (Katzner, 1988, p4), a representation of markets seems essential for the proper modelling of the future growth and evolution of cities.

Traditionally, economics (including urban economics) has focused on the equilibrium solution to a system. More recent research acknowledges the path-dependent nature of the growth of spatial economies and the bifurcations and other aspects of self-organizing and chaotic behaviour (Fujita *et al*, 1999). Such behaviour is often best analyzed with a simulation. Simulations of markets typically involve an abstraction of the price setting process, ascribing price movements to an imaginary "auctioneer" who either adjusts prices dynamically base on shortages or surpluses ("excess demand"), or experiments with pretend prices before establishing a set of market clearing prices (Katzner, 1988). This abstraction is not realistic (since few markets involve a real auctioneer and price adjusting, in reality, is done by the buyers and sellers, not a third party) but it has been accepted into economics because of its ability to simplify analysis.

The unrealism of the auctioneer is one limitation of traditional economics that has led to the blossoming of the field of Agent Based Computation Economics, defined as "the computational study of economies modeled as evolving systems of autonomous interacting agents" (Tesfatsion, 2002). Agent based computation economics allows one to abandon the traditional role of "externally imposed coordination devices such as fixed decision rules, common knowledge assumptions, representative agents, and market equilibrium constraints." (Tesfatsion, 2002) and allow face-to-face interactions to be modelled in more complex ways.

Direct simulation of markets involves abandoning the "auctioneer" and ascribing price movements to the actions of individual buyers and sellers. Most work in this area is aspatial. Seminal works include Gode and Sunder, 1993, and Steiglitz *et al*, 1995. In both of these works a "double auction" is adopted, where agents make offers to buy and/or sell, and a transaction occurs when a buy offer has a higher price than the lowest price sell offer.

In this work we adopt the following perspective:

- A long-range equilibrium model cannot handle the path-dependence and complexity that is inherent in real spatial systems, and so a time-series simulation is necessary.
- Within the simulation, it is doubtful that all of the complexity of geography, transportation, and heterogeneity of agents can be handled with algebraic models. Numerical modelling, perhaps including simulation, will be necessary.
- We are primarily interested in integrated models, that integrate the demand for transportation with a larger economic view of the households and firms that interact with each other in a region. Thus we must model supply-demand interaction.
- A market is a "mechanism facilitating exchange", and every interaction involves some exchange of something (be it goods or services, or mere pleasantries), so a flexible market representation could be used to represent any and all interactions.
- Agent based computational economics and disaggregate models need to be supported, but not necessarily embraced entirely.

Our objective is to describe a paradigm for thinking about integrating submodels based on the above perspective. The paradigm will suggest a software engineering framework, but it will also be shown that existing integrated models can be described according to the market paradigm even before they are adapted into a new software framework.

1.2 Continuous time discrete event simulation

As we have committed ourselves to time-series modelling, and the adoption of both legacy approaches and agent based computation approaches, we are interested in the simplest representation of time that supports both types of approaches. We assume that time is continuous, but that some things occur at instants. We adopt the notion of a global time clock that schedules the various algorithms in each of the submodels. A global time counter moves forward from the simulation start to the simulation end, and at any instant in time the various submodels can be called upon to simulate events.

A common implementation of the global time clock is the "event queue". The event queue consists of a list of events that are scheduled to occur in the future. The list is sorted by the time at which the events are to occur. Each event has a "handler" (or "event routine", see Lin and Fishwick, 1996) associated with it, which is an identifier for a software module and necessary parameters and code entry points. The global time clock selects the event that has the lowest time, and executes its handler, causing code to execute in the appropriate software module to handle the event. The handling of one event often causes future events to be added to the event queue. The simulation ends when the event queue is empty, when the time clock reaches a predetermined stopping time, or when the user terminates the simulation interactively.

A discrete event simulation can be executed on distributed computing facilities in parallel, by dividing submodels, agents or markets across processes on multiple processors. This requires either a "lookahead" value describing how far in time one process is allowed to get ahead of another process, or a "rollback" method to backup a process that gets too far ahead of other processes. (Lin and Fishwich, 1996).

1.3 Relationship to transaction simulation

This research follows from research by Abraham and Hunt, 2001, in building a transaction simulation framework. The transaction simulation framework represents full heterogeneity in agents, time and space. The notion of a market was adapted to handle all agent interaction. Agents make offers in markets at instants in time and at locations, and other agents may accept those offers at later instances in time, leading to a transaction identified at an arbitrary instant in time, at an arbitrary location, and between two arbitrary agents. Offers have prices associated with them, and prices change over time because individual agents adjust their offered price in response to their ability to meet their wants and needs. The transaction simulation framework is a method for a fully disaggregate representation of urban systems, with a behavioural representation for price movements and matching supply to demand. While research continues on demonstrating the effectiveness and long-term applicability of a transaction simulation paradigm, it has also become apparent that a whole-scale abandonment of all previous work using other paradigms is not likely to occur. Thus this paper shows how adopting the market paradigm supports ("embraces") the use of existing modelling techniques, while also facilitating the eventual extension of functioning model systems to more disaggregate, dynamic and continuous representations.

2. SIMULATION STRUCTURE

2.1 Offer list

A market is defined as an "institutional arrangement facilitating exchanges". We will make a minor assumption about how a market works to enable us to move forward: that a market consists of a list of published offers to transact. Thus as we move through time there will always be a list of outstanding offers. Each offer will be to transact a certain good or service at a place or range of places, and most offers will include a price (or price-like variable) that agents adjust to match supply to demand.

Integrating submodels into this framework then requires deciding for each submodel:

- which goods or services it supplies or demands,
- whether it does so as supplier or demander or both,
- whether it makes offers or accepts offers, or both,
- when it makes offers or accepts offers,
- what conditions and prices it attaches to its offers, and
- what information regarding the market it needs to decide what conditions to attach to its offers

2.2 The Auctioneer

Most views of markets rely on a story involving an "auctioneer" who controls the market itself, by adjusting the price to match supply and demand. The bulk of economic literature assumes a single price per "commodity" in each "place". When dealing with the imaginary auctioneer we will also assume that the auctioneer is working with a single price, but we may assume that the single price represents an average of observed prices. A single price only follows if we assume homogeneity within the category of good or service within a place. An average price requires that we work with large enough aggregations of space (in "zones"), time (in "time steps") and commodity types to have more than one transaction at each of the price points. Thus, the use of the auctioneer concept is seen as limiting for dynamic spatial economic models, since it requires an aggregation scheme in both time and space, or requires assuming homogeneity.

If there is a short term equilibrium assumption in a certain market, then during each time step the auctioneer must find the set of prices at the places so that supply and demand are equal. This is a tight coupling of supply, demand, and market mechanism, and usually a single algorithm in a single computational submodel solves for the equilibrium prices.

If there is no equilibrium assumption, then the "auctioneer" can be implemented somewhat independently of the agents providing the supply or demand. The auctioneer can simply track the shortage or surplus (the excess demand) over time, and adjust the price in response to the excess demand. In this case the auctioneer would adjust the prices of the outstanding offers in the market, and the auctioneer may offer limited or limitless quantities of supply and/or demand at the current price for the next time step.

An auctioneer is not required if the price movements are ascribed to a particular behavioural agent, so the consideration of individual agents is seen as important to the adoption of full heterogeneity of time and space.

2.3 Events affecting markets

Since we have defined the simulation as a series of events in time, and since we have defined all interactions as occurring in markets, we are only interested in the events that affect markets. These can be classified into six standard event types and an additional non standardized event type:

- offer event
- unoffer event
- transaction event
- price-update event

- query offers event
- query completed transactions event
- short term equilibrium event (non-standardized)

Offer event: An offer event occurs when a submodel or agent publishes offers in the market. These can represent actual decisions that are made by agents at specific times to participate in the market as an offerer to sell or purchase goods or services, or a wholescale specification of available prices and/or quantities by an aggregate submodel. Offers can contain a pointer to an acceptance handling routine (a "callback"), so that the offerer will be notified when its offer is accepted.

Unoffer event: An unoffer event occurs when an offer that was previously made (but not yet accepted) is withdrawn.

Transaction event: A transaction event occurs when a submodel or agent accepts an offer that is listed in the market. For example, when a simulated individual goes shopping, she will be accepting a number of offers to sell goods and services. When a transaction event occurs the offer accepter takes some action. As well, the offerer may also log the acceptance and take some further action through the callback acceptance handling routine.

Price-update event: A price update event is associated with the auctioneer in a market for goods and services. The quantities purchased and sold since the last price-update event are compared to calculate the "excess demand". The change in the price is a function of the excess demand, tending to decrease prices when excess demand is too low, and increase prices when excess demand is too high. A price update event corresponds to an auctioneer adjusting the price for all current offers in the market, and is often followed by the auctioneer offering infinite quantities of future demand and/or supply for the period of time until the next price update event.

Query Offers event: This event causes a query to occur on the list of outstanding offers, finding the set of offers (or a random subset of the set of offers) that meet certain criteria. This is used by agents or submodels to determine a set of offers that could be accepted. This event does not affect the market directly, but in many cases a Query Offers event is followed by a Transaction event, where one of the offers returned by the query is accepted.

Query Completed Transactions event: Agents or submodels sometimes need to be able to observe and respond to market conditions. This method allows agents to inspect past history, to determine past averages or trends and use those averages or trends in their decisions to make offers or accept offers.

In addition to the standard six types of events, there may also be **short term equilibrium events** in some markets**.** A short term equilibrium event is associated with an imaginary auctioneer, but is also associated with agents – both suppliers and demanders. It represents a process whereby a price is set that causes the supply and demand in a submarket to be equal over a time period, and is usually associated with just one computational submodel. The equilibrium is necessarily a *short term* equilibrium, representing only a subset of the decisions that influence supply and demand, otherwise the equilibrium event would be a complete solution to the entire model system, and our basic assumption regarding the necessity of a time series simulation would be violated and this paper would not be applicable. The method for calculating a short term equilibrium set of prices is usually highly specialized, and so cannot be standardized in the same way the other six types of events are standardized.

All interactions between agents in the dynamic simulation occur through one of these event types. Thus the basic implementation of the market paradigm integration framework involves an implementation of the first six events types, plus special code to implement any short term equilibriums that are represented.

2.4 Derivative (abstract) markets – the Aggregator

Integrated policy analysis models are often concerned with small details in space and time (e.g. the success of local corner stores, intra-day traffic patterns) yet still hope to forecast the long-range future of a region. To preserve the full integration of decisions, the multitude of minor details are often "aggregated up" into a higher level, more abstract, price-like variable, where they can influence the more major decisions. For instance, the attributes of travel between two points by the various available modes at the various times-of-day can be aggregated into a single composite utility number representing the difficulty of travelling between the two points. A common procedure in transportation modelling for aggregating detailed utility data is to use the logit model in the algorithm for choosing between details, which then allows using the "log-sum" of the available alternatives as the composite utility.

This can fit into the market paradigm if one imagines an aggregator who combines the prices and options in the detailed markets, and offers the aggregate alternative in a more abstract market. This "aggregator" is not usually a real actor, but it is no more abstract than the "auctioneer" of the typical story of how markets work. It is simply a description of how detailed information can be summarized for the benefit of submodels or agents.

3. LAG REPRESENTATIONS

3.1 Direct behavioural representation of lags

True lags occur in a system because of the rate of response of individual actors. If actors respond more slowly to changes in their situation, then there will be more lag. Actors respond slowly for two main reasons:

- it takes time and resources to notice and evaluate new conditions, and
- agents may understand that there is noise in the system, and forecast the future conditions based on some average of past conditions.

In a true dynamic microsimulation, the representation of lags would involve a direct behavioural representation of these two effects. Agents might only evaluate, say, their motor vehicle holdings every 6 months, reflecting the fact that people rarely go out and buy a car unless they have thought about it for some time. Agents may make their decision regarding which grocery store to shop at based on the average travel time to grocery stores over the past 8 weeks, rather than based on yesterday's traffic conditions.

3.2 Abstract representation of lags – auctioneer stories

The direct behavioural representation of lags is an attractive modelling approach. Unfortunately it cannot be used for all types of submodels. When there is an "auctioneer" story regarding price setting we cannot use behavioural observation of the actual (non-existent) auctioneer to determine the speed of price response, but we can still ascribe certain time lags in the behaviour to the auctioneer.

There are two basic lag representations, representing the two auctioneer "stories". If the auctioneer story is setting a price to achieve a short-term equilibrium over a time step, then a submodel considers the elements of supply, demand and auctioneer simultaneously. The amount of lag can be controlled by adjusting 1) the length of the time step, and 2) what elements of demand and supply are included in the "short term" equilibrium. If the auctioneer story involves the auctioneer setting prices in response to calculated excess demand, then the amount of lag can be controlled by adjusting the amount of time over which the excess demand is calculated, and the rate at which the price changes in response to excess demand.

3.3 Abstract representation of lags – leapfrog equilibrium

If more than one type of commodity is subject to a short term equilibrium calculation, then it would be ideal if the short term equilibrium for both types of commodities could be calculated simultaneously. In practice separate submodels usually handle the two equilibriums, and each usually assumes that the prices established by the other are fixed. Thus there is a "leapfrog equilibrium" lag introduced by the joint response of each equilibrium calculation to the prices established by the other equilibrium calculation.

4. EXAMPLE SUBMODELS

A typical **traffic assignment** submodel establishes a short term equilibrium between route choice and link performance. This submodel takes a trip table describing the demand for travel between points in space, and maps that trip table to demand for specific links on a network, adjusting the performance of each link in response to congestion. In typical untolled roadway networks, there is no price *per se* but the link travel times are "price like signals", in that they serve as the signal that causes the supply of travel on a link to equal the demand for travel on that link. The traffic assignment algorithm occurs as a short term equilibrium event. The event updates the costs of travel. The updated cost of travel consists of the "skims", or the matrix of place-to-place travel time and cost by travel mode. These skims are usually offered as fixed prices for any amount of future travel, and left as open offers until the next run of the traffic assignment submodel.

An **activity scheduling** module typically takes zone-to-zone travel prices by time of day and mode as constant, and determines the activity-travel patterns of individual household members. These travel patterns reflect the demand for retail goods (for shop trips), the supply of labour (for work trips), and the demand for various other interactions. The activity scheduling algorithm accepts offers of zone-to-zone travel every time it assigns a particular trip to an individual. Activity scheduling algorithms are usually able to calculate some measure of the composite utility of travel from a particular home location – "accessibilities" representing, in aggregate, the ability of individuals to meet their wants and needs given their home location. An activity scheduling algorithm could act as an aggregator, calculating these accessibility utilities and offering them to any household that wants to locate in the location corresponding to the accessibility. When a household accepts such an offer of accessibility (perhaps during the execution of a home location choice algorithm), the activity scheduling algorithm then knows it has to schedule activities for that household.

A **spatial price equilibrium** for spatial markets, such as in MEPLAN (Hunt and Simmonds, 1993) or PECAS (Hunt and Abraham, 2003), establishes the supply and demand for goods and services, and the demand for floorspace, over a range of locations (zones) and for a period of time. This equilibrium calculation is not a long run general equilibrium, because floorspace is usually considered fixed. In PECAS the equilibrium is consistent with economists' concept of a short-run equilibrium, because profits are allowed to be non-zero. In MEPLAN the long run equations are used (profit = 0) but it is still not a true "long run" equilibrium because floorspace is considered fixed. There are two additional sources of lag in these models. First, there are often inertia terms added to these models, which can be interpreted behaviourally as a cost associated with relocating, or a preference to remain in the same location as in the last time period. The inertia term represents long term aspects of the demand for location, so that PECAS or MEPLAN can focus on a short term equilibrium. Second, these models usually assume that travel costs by zone-pair are fixed within each time period, and the submodel is run in a leapfrog equilibrium with a travel demand model that assumes that the locations of activities are fixed in any one time period.

5. EXAMPLE SYSTEM – SACRAMENTO MEPLAN MODEL

To see how this perspective can be applied to existing modelling systems, we apply it to both the Oregon statewide model and to the Sacramento MEPLAN model. The intention is to show that although the description of the paradigm may be new, current modelling methods fit into the paradigm. Neither of these models currently conform to a software implementation of the market paradigm, but this paper will show that they could.

The Sacramento MEPLAN model is the simpler of the two, so we will describe it first. It's submodels and markets are shown in Figure 1. It was developed at the University of Calgary (Abraham, and Hunt, 1999) with funding and assistance from the University of California at Davis. It is being further enhanced and used for policy analysis by the Sacramento Area Council of Governments (Abraham et al, 2003).

5.1 Submodels

The model consists of 5 submodels:

LUSA: Land Use Allocation LUSA establishes the location of activity and the economic flows between zones. It takes floorspace and quantities and location of "exogenous production" (exogenous production is primarily exporting activity and retired households) as fixed and performs a spatial input-output allocation, adjusting prices to allocate activity. For goods and services, prices are taken to be equal to businesses' input costs.

LD: Land Development LD uses a logit model to allocate land quantities amongst future development states. The logit model utility functions are dependent on the current state of the property, the zoning regulations applicable to the property (so the model runs once for each current state for each zoning category, for each zone) and the prices established in LUSA.

FREDA/DERF: Trip generation and disutility calculation FREDA takes the economic flows from LUSA and translates them into trips. This module also takes the travel attributes calculated by TASA and calculates the economic costs of shipping or travelling by all modes, although when this calculation is being performed the module is called DERF.

TASA: Transport assignment TASA takes the trips from FREDA and assigns them to modes and routes. TASA summarizes travel costs by trip type and zone pair.

LUSB: Incremental allocation: LUSB takes the prices and disutilities established by LUSA and allocates new quantities of exogenous production to zones.

5.2 Markets

The connections between these submodels can be viewed in the market paradigm:

The **Link Travel** market: TASA performs a short-run equilibrium in the link travel market. The supply of travel links is taken as fixed in the short term, and the demand for travel is given as a matrix of trips between zone pairs. TASA allocates trips to routes and travel modes, adjusting the travel times on each link until a Wardrop user equilibrium is established.

The **Intrazonal Trip** market: At the end of a TASA run, TASA offers a matrix of travel costs by zone (“skims”). The quantity of transport supply offered by TASA is infinite, and the "prices" are in terms of cost and time (weighted averages by mode shares) and composite utility (the log-sum term) from the mode choice model. FREDA takes the economic flows associated with its offers that were accepted in the intrazonal shipping market (described below), and converts them to trips by accepting intrazonal trip offers.

The **Intrazonal Shipping** market: FREDA/DERF takes the offers made by TASA in the intrazonal trip market, and converts them to offers to ship quantities of goods and services. (For instance, for shopping TASA may offer specific times and costs for shopping trips. These offers are aggregated to offers to "acquire a dollar value of retail goods from a location" by DERF.) These are taken as offers at fixed prices and infinite quantities, and LUSA accepts these offers when it calculates its equilibrium.

The **Goods and services** markets: The markets for goods and services (the consumption of production quantities) are handled internally in LUSA. Each type of industry and household has elastic demand and supply coefficients, so LUSA contains the demand schedule, the supply schedule, and the price information. LUSA calculates a short-run equilibrium for goods and services.

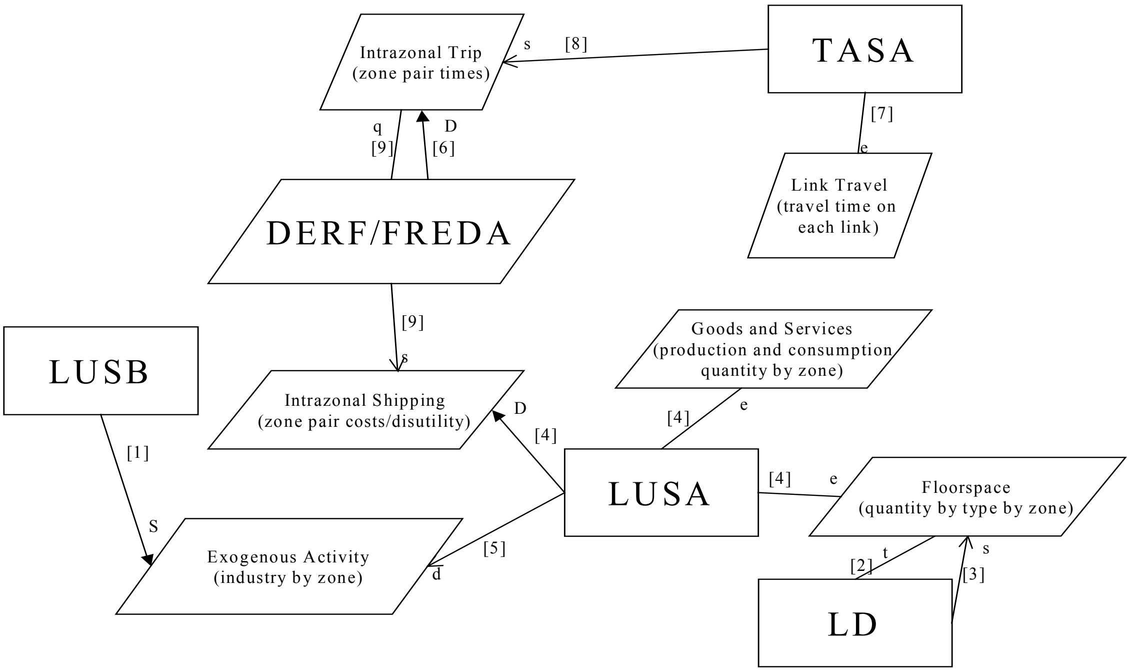

s = offer to sell/supply (supply)
d = offer to buy/demand (demand)
S = sell/supply (accept offer to buy) - also implies query of offers
D = buy/demand (accept offer to sell) - also implies query of offers
q = query the list of offers (without immediately accepting some)
t = query the list of completed transactions
p = price update
——▶ accept (transaction event, S or D)
——> offer (offer event, s or d) with callback
------> offer (offer event, s or d) without callback
[1] sequence of events

Figure 1 : Markets and submodels in the Sacramento MEPLAN model

The **Floorspace** markets: LD models developers' long-term behaviour, and offers the short-term quantity of floorspace. LUSA's equilibirum calculations consider the short-term demand curve for the floorspace in each zone, and establish the short-term equilibrium price for each floorspace type in each zone. LD considers these prices in the next iteration.

The **Growth in Exogenous Activity** market: LUSA establishes a set of prices and business conditions, and publishes these as offers for future exogenous locators. In the next time step LUSB allocates a total amount of growth in exogenous production to the zones according to the amounts published by LUSA.

5.3 Timing

The MEPLAN model in Sacramento runs in 5 year time steps. Within one time step first the LUSB module runs, changing the location of exogenous production by accepting the offers for locations of exogenous production. Then LD runs, querying the prices of floorspace established by LUSA, and simulating the actions of developers in producing more floorspace. The full amount of floorspace is then offered in the floorspace markets. LUSA then runs, calculating a short run equilibrium in the floorspace markets and the goods and services markets, accepting offers of intrazonal shipping. LUSA also updates the offers to locate exogenous production. FREDA then converts the accepted intrazonal shipments into intrazonal trips, accepting trip offers. TASA converts the accepted intrazonal trips into link travel, calculating a short run equilibrium in link travel, calculates the intrazonal trip travel times and costs, and updates its offers in the intrazonal trip market. DERF examines the trip costs in the intrazonal trip market, and updates its offers in the intrazonal shipment market. At the end of this process 5 years has elapsed (normally, the full 5 years is assumed to occur "during" the run of the LUSB and LD submodels, and the remaining submodels are assumed to occur in quick succession at the end of the 5 year period.

6. EXAMPLE SYSTEM – OREGON STATEWIDE MODEL.

The Oregon Statewide model is being developed by the Oregon Department of Transportation for policy analysis. It is described in Hunt *et al*, 2001. Again we describe the model in a way that is true to the theoretical design and mathematical implementation of the actual model, but as if the software implementation was consistent with the market paradigm. The submodels and markets of the Oregon Statewide model are shown in Figure 2.

6.1 Submodels

The model consists of seven modules:

ED: Regional Economics and Demographics The ED module provides the rest of the model with regional control total for production by economic sector, inputs and exports by economic sector, employment by labour category, in migration, and payroll by sector for each year.

PI: Production Allocations and Interactions The PI module determines for each year the distribution of production activity (industry and employment) among zones, the consumption of space by production activity, the flows of goods and services and labour from the location of production to the location of consumption, and the short term equilibrium prices for goods, services, labour and space. An aggregate allocation framework is used which allocates total production to zones, allocates technology within an industry in a zone, and allocates inputs and outputs in space. The system solves for the set of prices in space which clear each local (zonal) submarket for each category of good, service, labour and floorspace. The PI module is an implementation of the Activity Allocation submodel of PECAS (Hunt and Abraham, 2003), but with residential location choice by households removed.

HA: Household Allocation The HA module uses a microsimulation of the behaviour of individual households to determine changes in household composition, household actions regarding home location (and residential space use), final demand by households, household car ownership, employment status, employment location and career choice, and school status, school location and education choice. The price of residential floorspace is updated in HA based on the amount of vacancy.

LD: Land Development The LD module determines the actions of land owners and developers regarding the quantity and type of buildings. It uses a fine grid representation of the land and the associated regulations and characteristics regarding developability, and responds to the prices that are established for floorspace.

CT: Commercial Transport The CT module determines truck movements arising during a representative workday in the year. It takes the flow quantities by commodity type established in PI, and assigns them to modes and vehicles. Thus it converts economic flows to vehicle movements.

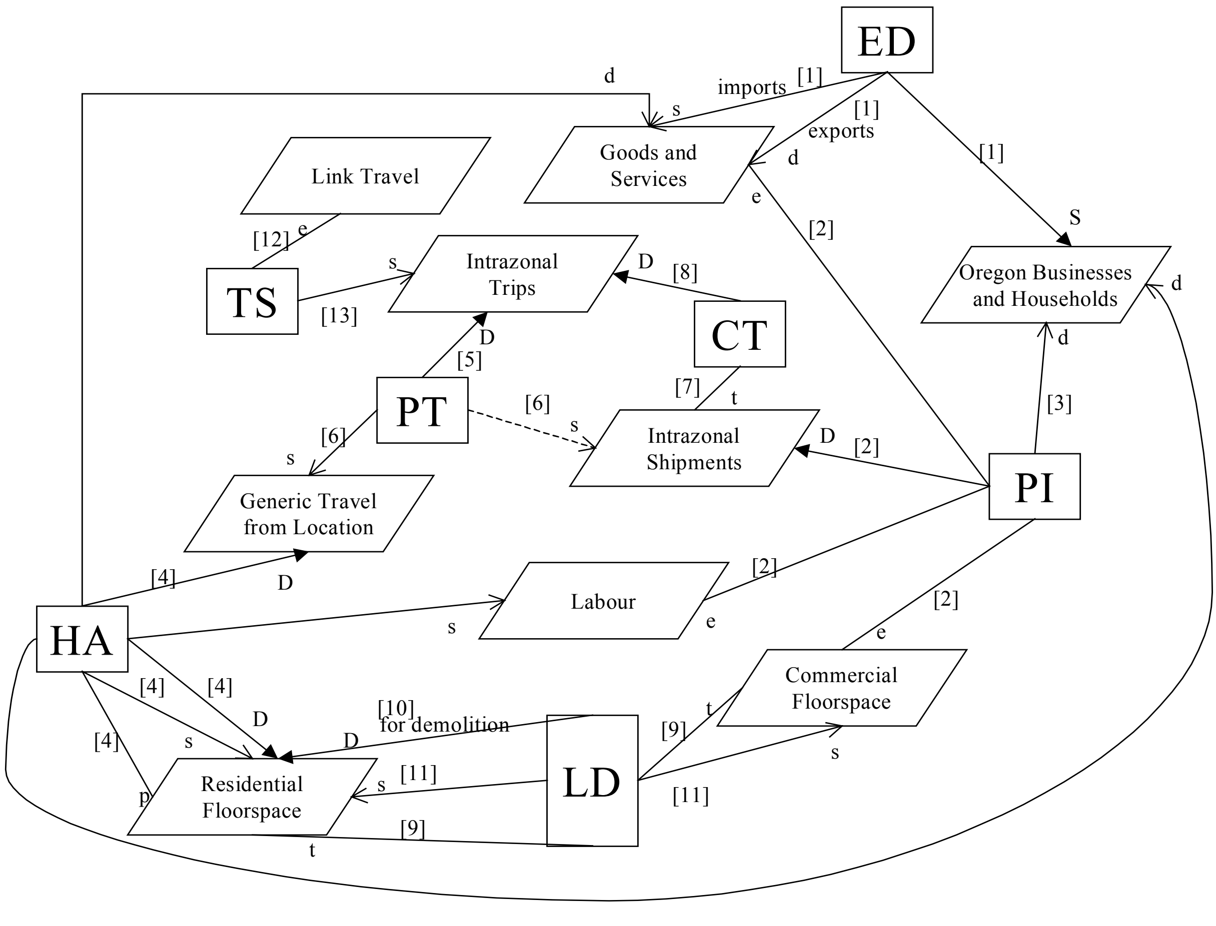

Figure 2 : Markets and modules in the Oregon Generation 2 Model

PT: Personal Travel The PT module establishes a list of the specific individual trips made by members of households during a representative day. It establishes for each trip a starting location, ending location, starting time, tour mode (mode choice) and vehicle occupancy. The process establishes an activity pattern for each person for the day (which is a listing of the sequence of activities undertaken by the household member), and then establishes the time of departure for each trip and the location for each stop on the trip.

TS: Transport Supply The TS module assigns the individual trips established in PT and CT to a network of transportation links. A shortest path algorithm is established for each trip, and after each trip's path is determined the performance of the links used by that trip are updated to account for congestion. Each tripmaker's preferences regarding route attributes are randomly selected. The algorithm iterates to a converged state, so that individual trip makers are choosing optimal routes given their preferences. The TS module reports link performance and zone-to-zone travel times and costs for a representative traveller.

6.2 Markets

The connections between these submodels can be viewed in the market paradigm:

The **Link Travel** market: This is the market for the physical space on a transportation link. In the short term the capacity is fixed. The Transport Supply module adjusts the routes of the trips in its trip list so that each traveller is choosing an optimal (for them) route, given the congested travel conditions. This is a short-term equilibrium event, and the TS module represents the behaviour of the supply, the demand, and the imaginary process (e.g. the "auctioneer") for establishing the set of link demands and the price-like signals of the travel time for each link.

The trip list reflects the short term demand, the fixed network reflects the short term supply. The long term supply is adjusted by the model user; the trip list comes from the Interzonal Vehicle Trip market.

The **Interzonal Vehicle Trip** market: This is the market for vehicle trips between pairs of locations. The Transport Supply module determines the travel distance and time by different modes and offers them as a set of matrices of representative times and costs for origin-destination pairs, using the zone system to delineate the origins and destinations. Thus TS offers an infinite quantity of travel between zone pairs, and publishes the money cost and travel times for a typical traveller for a particular mode at a particular time of day. Two other modules accept these offers: Personal Travel and Commercial Transport. PT and CT decide on the choices of where and when to travel for persons or shipments. TS represents the supply, and PT and CT represent different facets of demand.

The **Intrazonal Shipments/Flows** market: This is a derivative market, an aggregation and summary of the Intrazonal Vehicle Trip market. The Intrazonal Vehicle Trip market is for particular modes and at particular times of day. The Production Allocations and Interactions (PI) module does not work in terms of specific times of day or modes, but instead requires composite measures of the deterrence function for shipping between origin and destination. The PT module acts as an aggregator, calculating the expected maximum value of the mode and time-of-day decision, and offers the "log sum" value. The log-sum is a price-like signal, and it is offered in infinite quantity for a one year period. PI accepts those offers, leading to a matrix of interzonal flows that have not been assigned a mode or time-of-day. Note that although PT makes the offers, CT is responsible for generating the vehicle trips associated with accepting these offers. Thus in Figure 2 we can see that there is no callback needed to tell PT that its offers were accepted, but there is a "transaction query" event required so that CT can retrieve the list of offers that were accepted.

The **Generic Travel From Location** market: This market is an even more abstract view of travel. It represents the type of travel that households of a particular category make given a particular home location. The Personal Travel module PT again acts as an aggregator, calculating an accessibility value for different household categories as a price-like signal for household travel. These are offered in infinite quantity for a 1 year period. The HA module accepts those offers when it locates a particular household in a particular location. When HA accepts such an offer the callback routine could tell PT where a household has located. (In the existing Oregon implementation, with HA and PT running separately by year, the full list of households and their location is passed from HA to PT so there is no need for a callback.)

The **Residential Floorspace** market: The price of residential floorspace in each zone is established in HA through a price update function, which compares the vacancy rate in a zone with an equilibrium vacancy rate, and adjusts the price of all offers-to-let upwards if the vacancy rate is too low and downwards if it is too high. This allows a finer time-resolution in the response of prices – currently prices are updated monthly. All offers in this market are thus "market price" offers, in that there are no prices set by the offerers. The Land Development module offers newly constructed floorspace for rent. This new construction comes on once per year when the LD module is run. LD simulates developers decisions, and developers in this model make decisions based on the set of current prices, effectively assuming that future prices are likely to equal current prices. LD measures current prices by querying the transactions that have occurred (the 't' event in Figure 2). The Household Allocation module accepts offers of floorspace when a household chooses a new location. If the household is vacating an old location, that location is offered to the market, and becomes vacant. To accomodate demolition and reconstruction of occupied properties there is a continuous offer by each tenant to give up their property and supply it for demolition, and these offers are accepted by land-owners when they begin demolition.

The **Labour** market: The labour market is segregated by zone. The Production Allocation and Interactions module (PI) establishes a short term equilibrium of demand and supply of labour in each workplace location, and establishes a price for that labor. The demand for labour is calculated within PI directly. Since this is a short-term equilibrium calculation, PI must also have a representation of the supply of labour in each zone. PI has it's own simplified model of workplace destination choice for households to establish the supply schedule in each zone. (An alternative approach would be to have the price of labour in each zone established through a price update procedure. HA and PI could be price takers in supplying and demanding labour, and the price would be updated in response to the difference between the supply provided by HA and the demand consumed by PI.)

The **Oregon Businesses and Households** markets: PI and HA report on overall economic conditions in Oregon, and provide that input to ED for it's decisions regarding in migration. This can be put into the market paradigm by saying that PI and HA make fixed-price (actually, fixed utility) infinite quantity offers to move new households and new business activity into the model region. ED will accept a certain quantity of those offers.

The **Commercial Floorspace** markets: LD offers a fixed quantity of commercial floorspace in each time period. PI establishes a short term equilibrium with demand, accepting the offers of commercial floorspace, and establishing the price in each zone. These transaction prices are queried by LD to aid in developer decision making, with developers assuming that future prices will be the same as past year prices.

The **Goods and Services** markets: HA establishes the location and number of households, and hence establishes the final demand for goods and services by zone. These are taken as offers to purchase goods and services. ED establishes the import and export of goods and services from the model region: these are taken as offers to purchase and/or sell goods and services by agents exogenous to the simulation. These lists of offers are then all accepted by PI when it establishes the spatial distribution of goods flows. Intermediate flows of goods and services between businesses in Oregon are established within this PI short term equilibrium.

6.3 Timing

The model runs in one year time steps, in sequence through ED, PI, HA, CT, PT, LD and TS.

ED accepts offers to migrate businesses and households into the model region, and updates the offers to purchase and sell goods by importers and exporters.

PI then solves the short term equilibrium of business location, goods and services movements in abstract, labour, and commercial floorspace prices. In doing so it accepts household's offers to travel to work to supply labour and to purchase retail goods and services, landowners offers of commercial floorspace for rent, the offers to purchase and sell goods and services by importers and exporters, and the intrazonal shipment offers. PI then updates its own offer in the Oregon Business and Households market to accept new businesses in each industry category.

HA then moves households around, with households accepting offers for residential space, and adding new offers to let residential space corresponding to space vacated during a move. HA considers the accessibility numbers corresponding to offers of "generic travel given location" when choosing a location for each household. HA assigns particular jobs to individuals in households based on the flows established in PI. During the course of the HA operation, the prices of residential floorspace are updated by 12 price-update events. Thus HA is viewed as taking the full 12 months to complete, with a price update at the end of each month.

At the end of the year PT executes, generating the travel associated with households and businesses, and accepting offers in the Intrazonal Vehicle Trip market. PT updates the offered prices in the Intrazonal Shipments market, and in the Generic Travel Given Location market.

CT then executes, taking the intrazonal shipments that were accepted by PI and converting them to intrazonal vehicle trips, and accepting intrazonal vehicle trip offers. LD then executes, querying the residential floorspace prices established by HA and the commercial floorspace prices established by PI, and simulating the actions of developers who assume those prices are a guide to future rent revenue.

LD updates the offers to lease space in the commercial floorspace market and the residential floorspace market. When some residential floorspace is demolished or redeveloped, LD takes up households' offers to vacate their premises, evicting them.

Finally, TS takes all of its offers of intrazonal vehicle travel that were accepted, and converts them into link travel by solving a short term equilibrium of congested assignment. TS then updates its intrazonal travel offers.

6.4 Potential changes

The Oregon model contains some good example of how the dynamic transaction simulation paradigm could be more strongly adapted to further decouple modules and improve performance.

TS, PT and HA each process items individually: TS processes trips, adding them and removing them from links. PT processes individuals, choosing trips for them. HA processes households. TS offers intrazonal travel attributes to PT, and PT offers accessibility to HA. These could be run in parallel, with:

- PT processing individuals in a household once HA signals the end of its processing of a household by accepting an accessibility offer published by PT, and

- TS processing each trip for an individual as PT accepts TS's offers of intrazonal vehicle trips.

PI currently searches for a set of prices for each good or service, and each commercial floorspace type, in each zone, to clear each such market. PI could be made dynamic by replacing these price search processes with price update processes, with the same PI equations used to allocate consumption and production of goods and services and business location given a set of prices, but the set of prices change over time in response to shortages and surpluses, rather than to an equilibrium state in each time period that clears the market. This would make PI dynamic and would decrease computational burden, as the PI equations would not have to be repeatedly executed to establish an equilibrium (see also Abraham and Hunt, 2002).

7. CONCLUSIONS

This work developed a paradigm for *integrating* model components and agent representations of behaviour. Since integration involves interactions, since interactions involve some exchange, and since markets are mechanisms facilitating exchange, the mechanism is called a "market representation". This conveniently allows modelling to follow economic theory, and encourages the use of prices, or price-like variables, to match the quantity of demand to the quantity of supply and avoid (ongoing) shortages or surpluses.

The market representation is shown to be relatively simple, involving only six regular operations: make offer, withdraw offer, accept offer, query offers, query completed transactions and update market price. These six operations allow a decoupling of supply and demand elements in each market, and hence allow a decoupling of the various parts of a simulation and the various agents that interact to evolve cities and regions. Decoupling is desireable to facilitate plug-in code modularity and expansion, separate submodel development teams of analysts and programmers, and facilitate parallel execution on distributed computing facilities.

A seventh type of operation, a short term equilibrium event, requires a tighter coupling of supply, demand and market into a single equilibrium search algorithm. The concept of equilibrium does have certain advantages, and is so engrained into practice and economic research that supporting it seems essential. This paper shows how short term equilibrium events in some markets can be coupled with dynamic disequilibrium market representations for longer term trends in the same markets, and with fully dynamic representations in other markets. However this equilibrium event is a bit of a "force fit" into the paradigm – the true power of the paradigm emerges when equilibrium representations are replaced by fully dynamic representations.

Two existing modelling systems were investigated, to show how they fit into the market paradigm conceptually. Each of the individual submodels in this system could be wrapped up in interface code to make them fit into a software implementation of the market paradigm.

The market paradigm has a continuous time representation as its basis, and supports a continuous space representation. Existing modelling systems usually work in terms of time steps and aggregate space. It is disappointing to have to forgo the full potential of the market paradigm in this regard when taking on existing algorithms and submodels, but the alternative of throwing out most existing work and rebuilding from scratch requires too much of a risky commitment of research time and money to be appropriate for any planning agency, or even most research institutions. The work described here provides a forward path that may eventually lead to fully disaggregate agent based continuous time microsimulations, but that also has many safe and useful stops in the near future.

ACKNOWLEDGEMENTS

This work benefited from conversations with Eric Miller and Paul Salvini. The work was supported by the Social Science and Humanities Research Council of Canada under a Major Collaborative Research Initiative.

REFERENCES

Abraham, JE, GR Garry, AT Brownlee and JD Hunt, 2003, "Land Use Transport Modelling in Sacramento: Small Steps Towards A Grand Design", proceedings of the 8th International Conference on Computers in Urban Planning and Urban Management (CUPUM), Sendai, Japan, May 27-29, 2003

Abraham, J.E. and J.D. Hunt, 1999, "Policy Analysis using the Sacramento MEPLAN Land Use Transportation Interaction Model", *Transportation Research Record* 1685:199-208 and presented at the 1999 annual meeting of the *Transportation Research Board.*

Abraham, J.E. and J.D. Hunt, 2001, "Dynamic Microsimulation of Heterogeneous Spatial Markets", presented at the Workshop on Economics with Heterogeneous Interacting Agents, University of Maastricht, The Netherlands, June 7-9 2001

Abraham, J.E. and J.D. Hunt, 2002, "Spatial market representations: concepts and application to integrated planning models", presented at the *49th annual North American Meetings of the Regional Science Association International* in San Juan, Puerto Rico, November 14-16, 2002

Hunt, J.D., R Donnelly, JE Abraham, C Batten, J Freedman, J Hicks, PJ Costinett and WJ Upton, 2001, Design of a Statewide Land Use Transport Interaction Model for Oregon, World Conference on Transportation Research, Seoul, South Korea

Hunt, J.D and J.E. Abraham, 2003, Design and Application of the PECAS Land Use Modelling System, proceedings of the 8th International Conference on Computers in Urban Planning and Urban Management (CUPUM), Sendai, Japan, May 27-29, 2003

Katzner, D.W., 1988, *Walrasian Microeconomics*, Addison-Wesley, Reading, MA

Fujita, M, P. Krugman and A. Venables, *The Spatial Economy: Cities, Regions and International Trade*, MIT Press, Cambridge MA USA, 1999

Gode, D.K. and S. Sunder, 1993, Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality, *Journal of Political Economy*, 101:119-131

Lin, Y and PA Fishwick, 1996, Asynchronous Parallel Discrete Event Simulation, *IEEE Transactions on Systems, Man and Cybernetics,* Vol 26 No. 4 pp 397-412

Steiglitz, K, M.L. Honig and L. M. Cohen, 1995, A computational market model based on individual action. In Scott H. Clearwater, editor, Market-Based Control: A Paradigm for Distributed Resource Allocation. World Scientific, Hong Kong, pp1-27

Tesfatsion, L, 2002, "Agent-Based Computational Economics", Iowa State University Economics Working Paper No. 1, Revised July 2002, accessed at http://www.econ.iastate.edu/tesfatsi/acewp1.pdf on Dec. 19, 2002.