


STEM User Group Meeting 2010

Robin Bailey – Managing Director, ‘Mr STEM’



This space intentionally left blank

Day one – Tuesday 05 October

- 09:30 Showcase: modelling parallel technologies in telecoms networks and commercial or urban property developments
- 11:00 Capex work-out: fixed assets and depreciation, and a timing review for cashflow, NPV and equipment replacement
- 12:30 LUNCH
- 14:00 Platform evolution: priorities such as smarter results browsing driven by recent project experience
- 15:45: New market strategy: a small announcement to consulting partners with a potentially major impact on future growth
- 17:00 CLOSE
- 19:00 Reception: Saltmarsh Rooms, King's College
- 19:30 Gala dinner: Saltmarsh Rooms, King's College


Day two – Wednesday 06 October

- 09:00 Showcase: a jaw-dropping technique, and some personal recommendations on linking style
- 11:00 Opex work-out: consistent planning of maintenance, support staff and overhead costs
- 12:30 LUNCH
- 14:00 Technical spotlight: an end to confusion over incremental modelling of activity costs
- 15:45 Modelling clinic: live Q&A on current user roadblocks and some recent case histories
- 17:00 CLOSE
- 19:00 Informal meal for overnight guests

Keynote: imagine a world without network operators ...

- Visualise a company with offices in several different cities, or even countries, trying to establish a mechanism for exchanging messages and data between those offices
- Or consider a wealthy family, trying to emulate a similar system to stay in touch with cousins abroad
- The utilisation and geographical overhead of such an endeavour would be exorbitant (wouldn't it?)
- But there would be near universal demand for such a *service*
- The children of the 21st century know digital TV and radio, 3D cinema and two-way interactive media, and are more likely to use a laptop for gaming or Facebook than to develop software
- Programming and network design may have had their popular heyday
- The majority simply expect a network system to function, with little care or consideration for how it works
- There is modern gold rush towards systems hosting ...
- ... but there is really nothing new in the economic value of offering a slice of a large distributed system ...
- ... to a broad market of smaller players who could not afford to own such a system in their own right
- With few barriers to competition, many automation-based services will lose their value in the face of inevitable commoditisation and price erosion ...
- ... compared to the inescapable logic of providing a *network as a service*

An obvious economy of scale (think: millions of subscribers)




DIY ...




... or as a service

The cost of a truly private network


- There are two broad categories of communications transport:
 - fixed, encompassing everything (including string) from copper pairs and coaxial through to modern fibre; and
 - wireless, including point-to-point, 2G/3G/4G cellular, broadcast UHF and satellite technologies
- Total distance and cost are minimised by planning around a core network*



- For high resilience, a network must be planned with redundant paths so that no one element or link is critical



- A city may be more efficiently served with a single ring structure than forcing the traffic to a central core ...
- ... but clusters of sites are still better connected by a common core network



* Network planning is rather like transport planning: motorways are never built to individual addresses!

Getting a fixed network connection

- The fastest and only scalable contender for a modern fixed network connection is fibre, but establishing your own private fibre connection would be a costly endeavour
- The cheapest path is 'aerial'; i.e., suspended from poles every 35m:
 - imagine the visual impact of multiple private fibre networks!
- Underground trenching and ducts, with pits every 300m, is more expensive, and dependent on municipal planning
- Large carriers have built trunk fibres alongside railways or highways, or even along waterways:
 - it would be a planning nightmare if everyone wanted to run a private fibre along such sparse resources!
- Whereas a passive GPON network is designed for urban connections, these longer routes would also be served by active optical transports with repeater amplification every few 100km
- Where cluster of sites are aggregated onto a core DWDM system, Ethernet technology may be used to aggregate traffic from sites on one local ring to another
- The geographical extent and cost of maintaining such an end-to-end deployment is key to the economic incentive to share such assets
- If you had to supervise the delivery of every packet in a postal network first hand, you would become a full-time postman and never do anything else!


Getting a wireless network connection, and on the move

- Depending on spectrum, a point-to-point radio solution could be the most practical solution for a private network:
 - this has been the traditional solution for connecting regional TV broadcast towers
 - because of the point-to-point beam design of the radio signal, this technology is relatively immune to physical eavesdropping
- Very remote sites might be more effectively connected by satellite:
 - but at what cost (with backup)?!
- Multiple sites within a city might benefit from a technology like WiMAX:
 - but this is wide-open to physical eavesdropping, so
 - encryption is needed for 'privacy'
- If your employees are travelling, then none of these solutions are much good
- Providing a cellular radio system, in case one employee might pass through a given vicinity, would be as expensive as building private networks to every base-station site in the country
- Ironically, it would be even more expensive to roll-out a network of private phone kiosks, if they were present in every small town:
 - which is why the pay-phone has become much less prevalent since the advent of the mobile phone
 - in some remote regions, pay-phone services are provided from a GSM handset bolted to the wall of a container, to avoid the high cost of building a fixed connection

The pain and gain in distributed connectivity

- The good news is that, by creating a value chain of user and provider and offering the network as a service to users in many markets, an operator can share the cost of any common infrastructure between all of its users
- This is not just about using more of the ports on a network device, but also reducing the total distance overhead per user:
 - ultimately it is distance which is being 'delivered'
- The bad news is that, to reach this level of sharing and trust, the operator must offer his services to as many users in as many markets as possible:
 - and there are at least two harsh realities along the way


- The first matter is simple *utilisation*:
 - in a transport context, compare the cost of motoring for a single person with that of a family
 - a smaller car might cost half the price of a family car, but this is still 2.5 times more per head than for five riding in a family saloon
- The same applies to many network devices, such as routers and DSLAMs:
 - until you achieve critical mass, low utilisation of a device can equate



to a 5x overhead on cost per head

Unused-capacity overhead increases with geographical scale

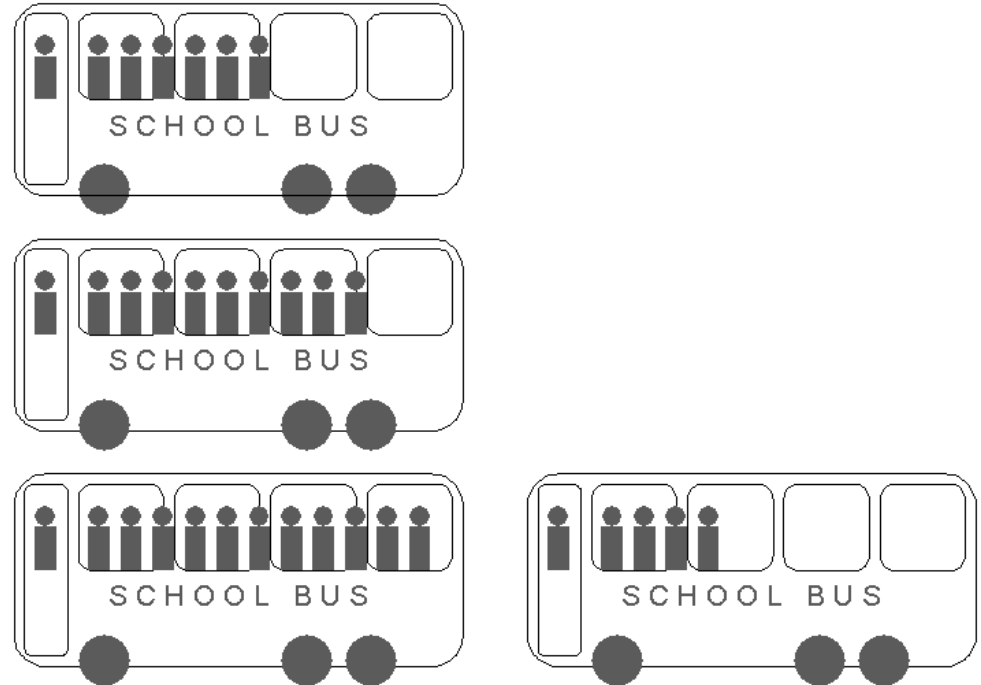
- The second is a slightly more subtle issue of *geographical deployment*, often overlooked in simplistic back-of-envelope calculations or hurried spreadsheet analyses
- If your target market encompasses many towns or cities, then it is very unlikely that your network will revolve around a single switch fabric
- Competitive pressure denies the luxury of achieving scale in one location before offering service in the next
- So the practical impact is that you have to sustain sub-scale utilisation in many geographical locations to start with ...
- ... and every location continues to add an overhead of slack (i.e., unused capacity) on an indefinite basis



- Mobile networks have to provide coverage, when there is low initial utilisation at all base-station sites...
- ... but natural variance in demand means that there will continue to be an overhead of unused capacity at each separate site

A transport analogy from the education sector


- Consider an education authority providing buses for children to travel to a number of schools in a region:
 - each individual bus might have a capacity of 50 pupils
 - but each school will require its own bus if everyone is to get to school on time
 - though there might only be 400 distant pupils across ten schools
- Worse, while some schools may have less than the average of 40 per school, there will be some with enough to need a second bus!
- Even with several bus loads per school, there still might be an average of half the seats empty on one bus for each separate school, due to averaging of numbers across the different schools



- These technical finance and economic issues are covered in depth in our STEM training course written around a model comparing *WiMAX* and *DSL* as *alternative technologies for the provision of broadband in rural areas*

Tapping into an always-on revenue stream

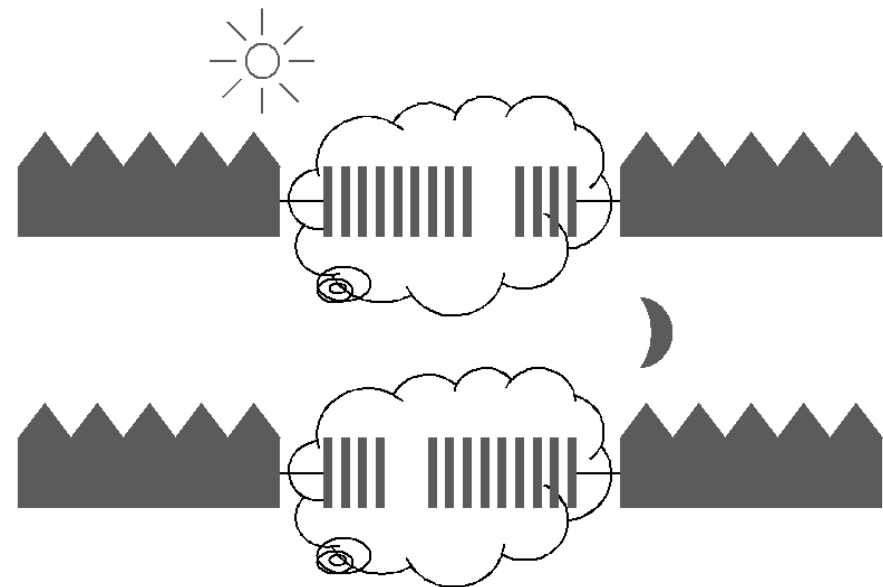
- Building a communications network outside of a lab. is a costly business ...
- ... but when the network is switched on and offered to fee-paying 'passengers', then it will hum with activity from day one
- Consider a modest 40 x 40Gbit/s λ DWDM system, which might cost €1m for the optical interfaces and electronics, €2m for a 1% share of a 1000km buried fibre, plus another €1m for repeater amplification along the way



- Suppose it is only 1% utilised on average over a 24-hour period
- This implies an average delivery of 16Gbit/s:
 - 16 billion bits per second
 - 960 billion bits per minute
 - 57,600 billion bits per hour
 - 1,382,400 billion bits per day
 - 504,576,000 billion bits p.a.
 - 63,072,000 billion bytes p.a.
 - 58,740,377 GB p.a
- At €0.16 per GB*, this yields a jaw-dropping annual revenue of €9.4m, compared to the annualised cost of €0.6m:
 - €8.8m profit at a margin of 94%

Reaping the benefits of service-provider expertise

- Some of the economic aspects of managing computing infrastructure look like the wider issues of maintaining a telecoms network in microcosm
- For smaller companies, the cost overhead of always-on high-bandwidth connectivity may be an expensive luxury if demand arrives in isolated bursts, whereas a hosting service enables the cost to be shared among many clients with similar requirements, yielding a lower unit cost for the client while affording a margin to the operator
- Small and medium enterprises also benefit from leasing options for fully-managed computing solutions if they do not have the scale to employ their own round-the-clock support staff
- Larger companies, and many Internet-based businesses, which may place wildly varying and unpredictable demands on their processing infrastructure will increasingly benefit from cloud computing services which provide virtualised processor solutions on-demand*



The principal value comes in sharing the infrastructure

- Proportionately less overhead is required when the processing power is all in one 'farm', compared to each client maintaining its own, separate computing facilities
- At higher utilisation, it is also practical to replace the equipment more frequently, so everyone benefits from the latest technology
- An on-demand solution enables rapid deployment of infrastructure, and moves the required investment into monthly payments (opex) instead of up-front capex (good for start-ups)
- The overhead of training IT staff in best-practice security, reliability and disaster-recovery techniques becomes the preserve of the service provider, allowing the client organisation to focus on its own business imperatives
- But what is the likely end-game as other players join the rush to offer cloud services?
- Competition is fierce, and will become only more so, given the low barriers to market entry *
- A market leader might innovate today, but the product is ripe for automation and will become rapidly commoditised
- And IT technology is still improving, in terms of ease of operation / reliability:
 - once too hard, but soon too easy?
- A small company can own and operate its own servers, even in a hosted environment, more economically than leasing in the cloud:
 - it may be happy to suffer a once-a-year server outage if it halves the IT bill all year round

Choosing between innovation and necessity

- Which elements of the value chain are vital (i.e., you can't afford the alternative) and which are just convenient?
- A hosting provider is providing network connectivity, round-the-clock monitoring, multi point-of-presence path-protection and resilience too
- These are all attributes of a service provider which are very expensive to operate on a private or dedicated basis
- A cloud provider is certainly enabling better utilisation of computing resources than will be achieved by a business with widely varying demands
- But many businesses with more predictable requirements will be better served by an optimised solution in the form of a private cloud infrastructure
- Retail service providers will stand to succeed solely on the strength of their brand and customer loyalty, as efficiency and standardised internal processes inevitably converge
- In contrast, the distance 'delivered' by the network operator is an economic absolute of *indispensable connectivity*
- Market forces will constrain the margin which may be earned from a network, but it will be an evergreen investment
- More lucrative services based on IT innovation will ride an economic seesaw of competition and consolidation as technology evolves and market requirements change
- A balanced investment portfolio will be wise, given the near certainty that our desire to communicate and transport data will only ever increase

STOP PRESS

- Unless, of course, there is a market-changing event ...
- ... such as the discovery of a new ultra-long wavelength radio technology which could combine high-bandwidth information rates with the same capacity for long-distance communication enjoyed by whales!
- Then it might be practical to carry about with you your own private network interface, and the age of the *network as a service* would be over




- But I am not holding my breath ...



This space intentionally left blank

Showcase: modelling parallel technologies in telecoms networks and commercial property developments

Robin Bailey – Managing Director, ‘Mr STEM’





A change in name

- Look carefully at the two STEM 'about boxes' shown alongside
- Besides the evident company re-branding, the new version proclaims:

industry

- the STEM process is highly flexible and more widely applicable than 'just telecoms'
- its automation helps you get things done
- This subtle alteration marks a significant ambition for the new Implied Logic Limited ...
- ...which has a rather different mission to its forerunner in business! (More on this later.)

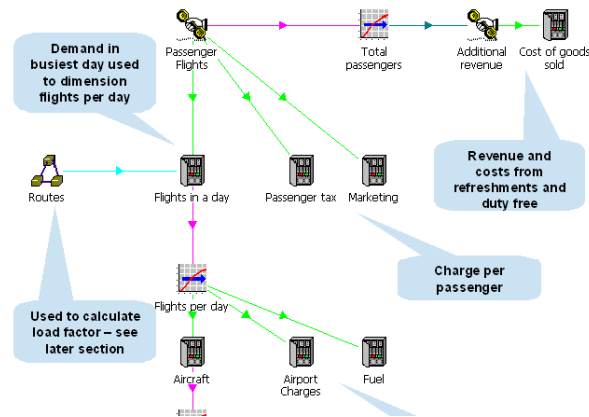


The STEM modelling process can be applied to a variety of sectors beyond the initial learning ground of telecoms

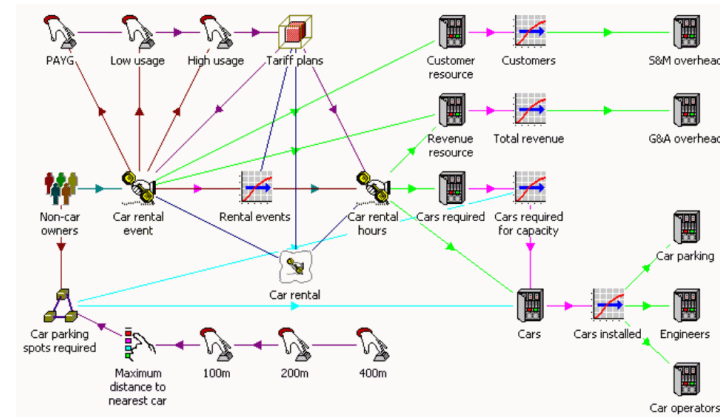
- In previous User Group Meetings (2007–9) we have used STEM to model:
 - budget airlines, car hire, energy generation, IT, and online conferencing
- In the last twelve months we have been involved in a diverse range of projects including:
 - parallel GPON, wireless and satellite access networks
 - layer 2 and layer 3 NGN aggregation
 - integrated IT services for commercial properties
- For reasons of client confidentiality, sadly it is not possible to share the specific details of these models
- Instead we will present a progression of idealised models – from classical network modelling, through cloud services and facilities for serviced offices, and on to services for commercial developments – which will serve as an eye-opener to the massive potential which remains largely untapped

Our experimentation has been public and interactive

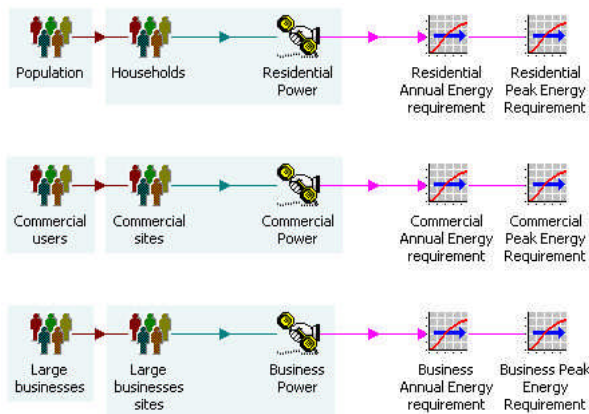
- Exploring strategic choices for a low-cost airline



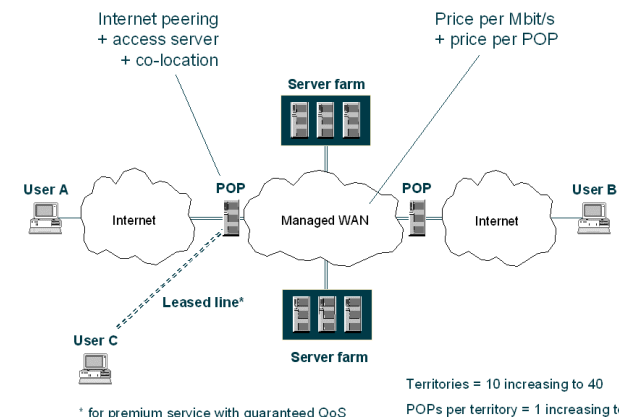
- Exploring a business case for distributed car hire



- Optimising the mix of base-load and peak energy generation



- Calculating the value of an online conferencing service provider



A progression from networks to information technology

■ Classical network modelling:

- fixed access
 - wireless access
 - satellite access
- } + aggregation + transport + switching

■ Cloud computing services:

- Hosting, virtual servers, computing on demand, security, backup
- Data centres, peering, racks, servers, storage, asynchronous replication

■ Facilities for serviced offices:

- Offices, virtual office addresses, itinerant office space, meeting rooms
- Commercial buildings, furniture, comms infrastructure, reception

■ ICT services for commercial property developments:

- Core network, building-management, security, RFID, advertising
- Separate modelling views for plan, build and operate phases

The key drivers for fixed and wireless networks

SERVICES

- Business and consumer:
 - voice
 - broadband (~2Mbits/s) and high-speed broadband (~20Mbits/s)
 - utility metering
 - TV
- Commercial:
 - wireless backhaul
 - commercial interconnect
- Regulated/government:
 - emergency 911
 - legal intercept
- Police/intelligence services/military:
 - secure communications



RESOURCES

- Fixed access: optical termination, final drop, fibre per subscriber, splicing, duct and trenching, FDH (splitters), optical frame and line termination
- Wireless access: CPE, spectrum, antennae, carriers, BTS, site lease, microwave links, RNC, backhaul
- Satellite access: CPE, transponders, uplink centre, interconnection
- Aggregation: Ethernet fan-out, 2x protection, 10G/100G ports, cards and shelves, links to DWDM
- Transport: link transponders (4x10G, 100G; paired), cards and shelves, DWDM wavelengths, fibre, repeaters
- Switching: dual homed, Ethernet aggregation, 10G/100G ports, cards and shelves, commercial interconnect

The key drivers for cloud computing services*

SERVICES

- Hosting:
 - website, applications, email, dedicated servers
 - varying service levels
- Virtual servers:
 - applications, including fileserver, database, email, website, intranet
 - varying specifications
- Computing on demand
- Firewall, security, backup and recovery
- Human interfaces:
 - commercial management (CFO)
 - technical management (CTO)
 - technical support (users)



RESOURCES

- Data centres:
 - internal network and switching
 - external peering
 - production racks
 - power, generation and cooling
 - servers
 - software licensing
 - storage
 - asynchronous replication
- Human resources:
 - Network architect(s)
 - delivery engineers and technical management
 - user support staff

The key drivers for serviced office facilities*

SERVICES

- Offices:
 - space, rooms, furniture and fittings, security
 - dedicated phone number, PABX lines, direct fax lines, shared and dedicated Internet bandwidth
 - visitor reception, switchboard, print and copy, messages, post
- Virtual office addresses:
 - dedicated number, switchboard, messages, post handling
- Itinerant office space:
 - walk-in PAYG access to a network of office suites / lounges
- Meeting rooms:
 - training facilities, video-conferencing, telepresence



RESOURCES

- Commercial office building or space (leased or owned), air-conditioning
- Furniture assets
- Utility charges and insurance
- Common network and comms infrastructure
- Common entrance, kitchen and washroom areas
- Per-client mail storage area
- Office managers, reception staff, cleaning, security
- Common health and safety
- Generic template for a serviced office replicated across multiple national/worldwide locations
- Comparable methodologies for hotels, restaurants, shopping malls, ...

The key drivers for commercial property ICT services*

SERVICES

- Core network connectivity: fibre, CAT3, wireless mesh; private networks, LAN, WLAN, Internet; voice, fax, PABX, videoconference, hosted email
- Integrated building-management: light, power, metering, backup generators, UPS, lifts, parking, fire detection and sprinkler systems, monitoring, health and safety
- Security: CCTV, access control, intruder alert, public address
- RFID: asset tracking, visitor location, retail security
- Cash and payment: ATM, POS
- Advertising: digital signage, interactive terminals
- Environs: signage, street lighting, drainage, energy distribution, disaster planning



RESOURCES


- Network assets (owned or leased)
- Buildings, facilities and systems assets
- Architect, commercial, planning, legal and financial teams or sub-contractors
- Management, delivery, operations, maintenance and audit teams
- Generic templates for different classes of buildings replicated across numerous construction sites
- Separate modelling views for plan, build and operate phases
- Scenarios for different construction dynamics and growth trajectories



This space intentionally left blank

Capex work-out: fixed assets and depreciation, plus timing for cashflow, NPV and equipment replacement

Robin Bailey – Managing Director, ‘Mr STEM’




Introduction

- Automating the calculation of capex and depreciation is the ‘bread and butter’ of a STEM model which saves users writing repetitive and rudimentary formulae and ensures consistency and reliability
- We will review the basic ingredients of this methodology:
 - relevant inputs and results
 - different options for depreciation
- And then examine the assumed timing for cashflow and NPV
- Finally we will examine some specific issues arising from clients:
 - capitalised installation costs
 - equipment replacement
 - replacement capex
 - initial capex

Relevant inputs and results

INPUTS

- [Unit capacity and] physical lifetime:
 - quantity and timing depends on incremental demand for the asset*
- Unit capital cost, plus optional:
 - residual value
 - tangible flag
- Calibration period, capital cost trends (local and global), capital cost structure
- Price list (economies of scale)
- Depreciation options:
 - financial lifetime
 - depreciation rate
 - depreciation schedule
 - calendar override





RESULTS

- Installed and Incremental Units results multiply into per unit assumptions
- Capital Expenditure:
 - tangible and intangible
- Depreciation and amortisation
- Proceeds from Sale of Assets:
 - the positive cash flow item that arises when the Resource is sold
- Profit on Sale of Assets:
 - arises if a Resource is sold for more than its net book value
 - normally zero, but could be positive if the Resource is removed early and is sold for a higher price than the original estimated residual value

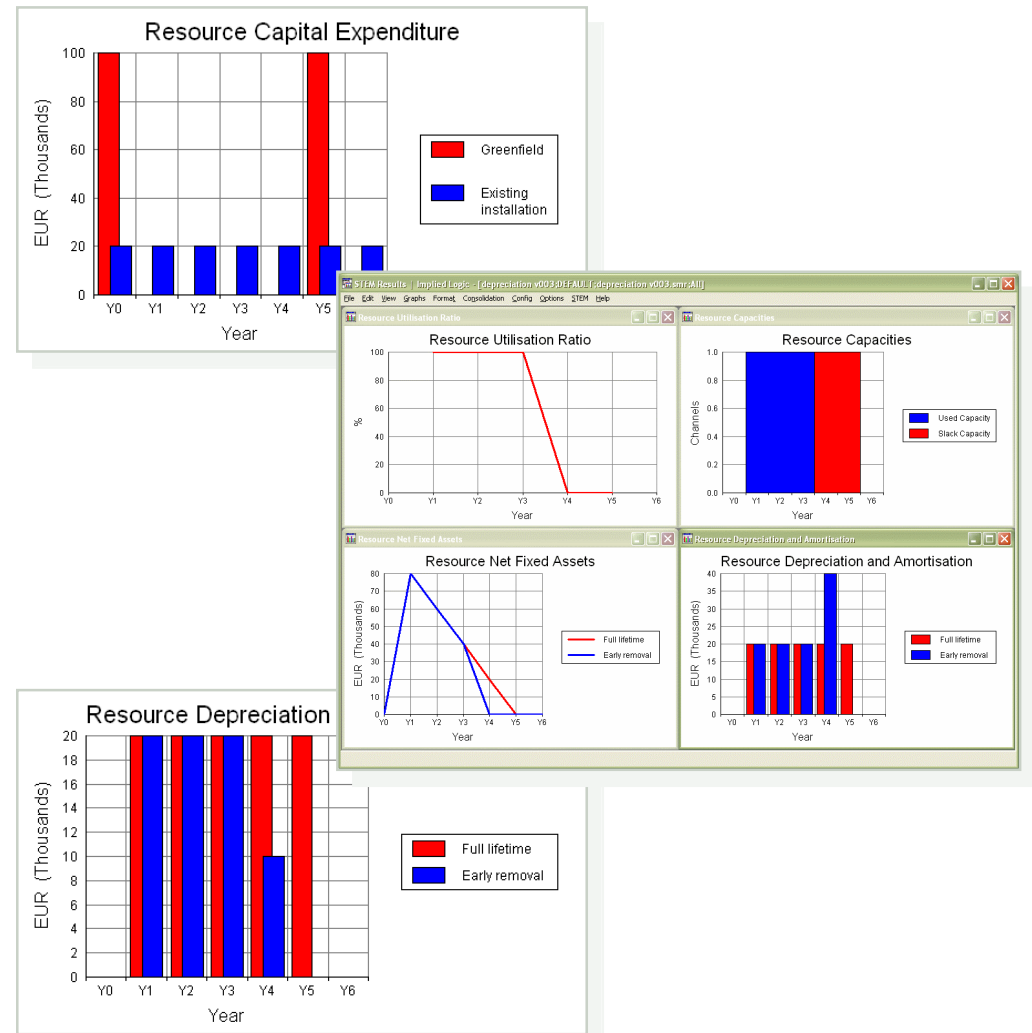
Different calculation methodologies

- **Straight-line**: depreciation is charged in equal portions over the full [physical] lifetime of a Resource (by default),
 - preferred for service costing as it reflects the business value of the asset, assuming used throughout
 - a shorter financial lifetime is typical for accounting purposes, allowing for possible obsolescence
- **Reducing balance**: charges a certain proportion of written-down value each year, accelerating the write down
 - MACRS in the US averages the last two years' depreciation charges
- **Depreciation schedule**: roll your own!
- **Calendar override**: allows you to model a re-planning exercise where depreciation must be accelerated in anticipation of early replacement



Refinements for initial deployment and early removal

- Pre-run installation:
 - avoids all capex appearing in Y0 if assets were pre-existing
- Redundant unit write-off:
 - allows a proportion of no-longer used equipment to be removed from the business and sold on
 - any remaining written-down value must be charged as depreciation
- Residual value:
 - sell-on value may be different if withdrawn earlier than expected*
 - straight-line depreciation actually rarely over-stated when originally calculated for the expected full physical lifetime






* Example shown has zero value after full lifetime, but sells for 30% two years early

Timing of capex, depreciation and re-sale

EUR	Y0	Q1	Q2	Q3	Q4	Q1 Y2	Q2 Y2	Q3 Y2	Q4 Y2	Q1 Y3	Q2 Y3	Q3 Y3	Q4 Y3	Q1 Y4	Q2 Y4	Q3 Y4	Q4 Y4	Q1 Y5	Q2 Y5	Q3 Y5	Q4 Y5	Q1 Y6	Q2 Y6	Q3 Y6	Q4 Y6
Capital Expenditure	0	100,000	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Depreciation	0	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	3,500	0	0	0	0
Proceeds from Sale of Assets	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	30,000

- If a resource installed for five years from Y1-Y5
- The capex is assumed to arise at the *start* of Y1:
 - or Q1 if the model is run in quarters
 - or Jan/M1 in months
- Depreciation is not an *event* as such:
 - it is an accounting charge for a period
 - calculated in proportion to the length
- Net book value is reported for the *period end*:
 - so the first value shown is never the full capex, but net of the initial depreciation*
- Any re-sale income (residual value) occurs at the *start* of the period after a resource is removed




* Unless there is a deferred depreciation schedule

What this means for borrowing, equity and NPV

- Borrowing is calculated at the end of each period, driven by the difference between retained earnings and net fixed assets
- It is easy to think this means that STEM will ‘under-borrow’ because it will never see the initial, full book value
- The requirement is actually reduced by:
 - by in-period revenues
 - deferred payment terms
- In the absence of any revenue or share capital, STEM will borrow to the full value:
 - the reduced net-asset value is balanced by the negative equity*
- NPV appears to under-state the loss too:
 - but depreciation is not a factor at all
 - the cashflow is simply discounted!

Balance Sheet

EUR	Y0	Y1	
Cash	0	0	
Cash Deposits	0	0	
Debtors	0	0	
Current Assets	0	0	
Net Tangible Assets	0	80,000	€
Net Intangible Assets	0	0	
Investments	0	0	
Long-Term Assets	0	80,000	€
Total Assets	0	80,000	€
Overdraft	0	0	
Long-Term Borrowing Payable	0	20,000	
Creditors	0	0	
Tax Payable	0	0	
Dividends Payable	0	0	
Current Liabilities	0	20,000	
Borrowing Payable after One Yr.	0	80,000	€
Total Liabilities	0	100,000	€
Share Capital	0	0	
Retained Earnings	0	-20,000	-€
Equity	0	-20,000	-€
Total Liabilities plus Equity	0	80,000	€



Issues arising: capitalised installation costs

- Resources have numerous built-in cost attributes, such as capital, operations and maintenance costs (per unit)
- The capital cost is the only attribute which is a capitalised, one-off cost
- But what if you need to capitalise some other aspect of a resource's cost, such as an associated installation activity?
- Although you can split-out different components of capital cost with user data (and define the actual capital cost input through a formula) ...
- ... the separation is only visible on the input side, and there is still only one Capital Expenditure result
- Using a separate, linked resource for the installation cost is evidently a tedious solution if the majority of resources must then be duplicated!
- When one customer proceeded along exactly this route, we decided to look harder at the problem
- It is not hard to create custom results which calculate individual capex and straight-line depreciation outputs for the components of an aggregate capital cost input ...
- ... at least if pre-run installation is not being used
- This is much more flexible for the user than a further built-in attribute
- Such a solution is usually implemented by our support team, as it requires knowledge of the internal syntax
- We will only add more built-in attributes if a more regular pattern of requests emerges:
 - “small is beautiful”

Issues arising: equipment replacement


- To avoid unnecessary complexity in calculating – and understanding – depreciation results, all resources installed in a given period will finally expire (i.e., come up for replacement) at the same time
- In typical scenarios where demand is growing, even slowly, this is often enough to effect a natural spread of replacement
- One new client sought a closer match with ‘reality’ in the shape of some kind of averaging effect
- We then considered implementing a ‘replacement profile’ input to prescribe a spread of actual lifetime:
 - similar format to Pre-Run Profile
- Actually partitioning the calculation of depreciation based on prior knowledge of such a spread would be unrealistic, as well as unmanageably complicated
- There is no problem with accounting depreciation over a shorter financial lifetime
- When allocating the value of an asset over ‘its full lifetime’, then we could :
 - have a ‘sting in the tail’ for early replacement, and an allocation holiday for extended lifetime, or
 - allocate bonus value from the overrun to compensate for under-reporting from the early expiry
- But it would be much easier to audit three similar resources in a function! *

Issues arising: replacement capex

- Consider the roll-out of a new, national network over a ten year period:
 - the elapsed time between starting in the most populous urban areas and reaching the most remote rural locations
 - subject to constraints of labour or capital commitment
- Whereas in a shorter project, you could regard all expenditure as new build ...
- ... on this longer timescale, some equipment installed early on will need to be replaced before the project end
- Therefore the total cumulative capex may misrepresent and overstate the initial budget for the development
- In a STEM model, resources are installed to meet incremental demand
- For each calculation period, STEM compares the existing capacity used by a service with the current demand from that service, and installs new equipment (or takes up existing slack capacity) if there is a shortfall
- So incremental units may arise from new demand, or replacement, or both
- Prior to STEM 7.2a, STEM did not distinguish these cases
- Calculating the initial capex required running a scenario where all lifetimes were artificially set in excess of the run period to avoid replacement altogether

Distinguishing initial and replacement capital expenditure

- STEM has always reported the Installed Units and Incremental Units for each period of a resource's lifetime
- STEM 7.2a store the number of Expired Units in any period as a new result which shows the number of units which have reached the end of their physical life (and may need to be replaced) in a given period
- An inferred Replacement Proportion for Incremental Units is then stored for each new installation 'age'
- This allows for a full separation between new and replacement for both capital expenditure and all associated depreciation



Making some assumptions about the replacement proportion

- The slightly tenuous part is working out what proportion (if any) of those expired units is being replaced
- If demand is flat or increasing, then an easy assumption is that all expired units are replaced ...
- ... while if demand ceases altogether, then there are no incremental units and evidently none are replaced
- But what if demand is reducing? Not all expired units need to be replaced ...
- ... but we make the assumption that, if there are any incremental units, then the lesser of incremental and expired units can be regarded as replacement
- So if there were three incremental units and two expired units, then we would regard two of the three new units as replacements
- Conversely, if there were two incremental units and three expired units, then we would regard both new units as replacements
- But if a new service re-uses existing equipment as a former service wanes:
 - is a new box installed by the new service a replacement for an old box used by the declining service?
- In the majority of cases, we can identify an unambiguous Replacement Proportion of Incremental Units in any calculation period

Initial and replacement capex/depreciation results

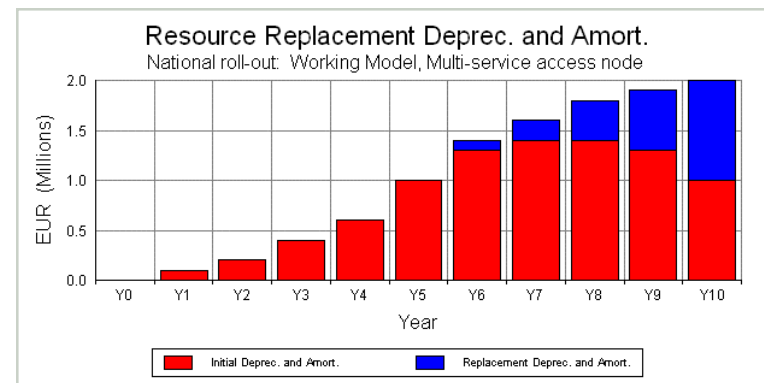

- In our specific example, demand is always increasing, so Incremental Units consistently exceeds Expired Units and we get the following results for the Replacement Proportion

Units	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Expired Units	0.00	0.00	0.00	0.00	0.00	0.00	50.00	50.00	100.00	100.00	200.00
Incremental Units	0.00	50.00	50.00	100.00	100.00	200.00	250.00	150.00	200.00	150.00	250.00
Replacement Proportion	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.33	0.50	0.67	0.80

- Based on this Replacement Proportion, STEM generates a Replacement Capital Expenditure result, and an Initial Capital Expenditure complement (also available for the network as a whole)

- And by remembering this proportion for each new installed age of a resource, STEM can also generate Initial and Replacement Depreciation results

- The symmetrical pattern of the initial capex below suggests a faithful classification of the split between ‘growth’ and ‘replacement’ capex



Issues arising: initial capex

- It might seem backwards to consider the 'initial capex' after replacement capex
- But this is actually an abbreviation for the rather vexed topic of initial capex to be considered in a model, *allowing for what is historically already on the books in real life*
- The specific challenge here is to match the ongoing depreciation calculation in STEM with actual accounting records
- This topic is better explained by a practitioner who has worked on this specific problem for a number of years

Issues arising: match depreciation calculation in STEM with accounting records


- Modelling of life networks with a “history” for different purposes:
 - costing
 - replacement / migration / network extension / what if scenarios
- Accounting records (asset register) normally don't match model structure, which has been developed for the special purpose of the exercise:
 - resources in the model are chosen to reflect technical planning processes
 - dedicated and shared resources for costing
 - transparent modeling of roll-out strategies
- Assets have been re-evaluated, special write offs have been undertaken during the lifetime of the network
- Equipment cost has changed over time (quite different per asset type)



It's quite hard to use pre-run data (probably even backwards engineered) to match data in accounting records

Solution: match depreciation calculation in STEM with accounting records (1)

- Use 'standard' depreciation calculation in STEM for new assets
- Use pre-run installation to have the capacity available at model start ...
- ... but avoid depreciation calculation for those resources (via a cost trend)



assets - Cost Index "Switch off capex"/Cos

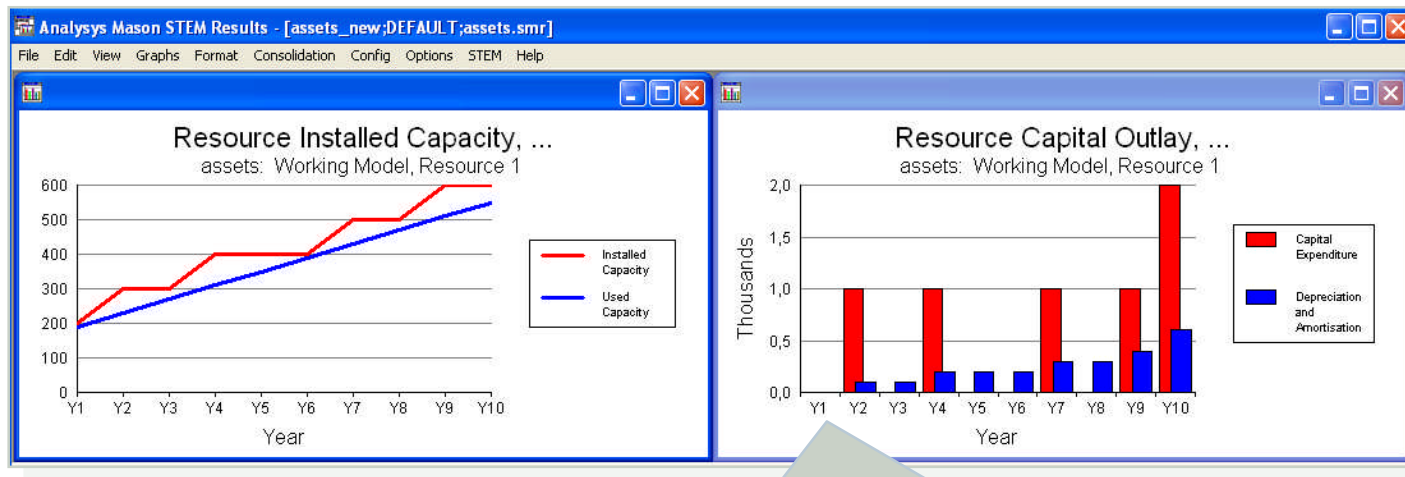
Close Edit Variants Move Type

X ✓ 0,00

Period Y0 Y1

Value 0,00 1,00

Capacity and Lifetime	
Capacity	
Unit Capacity	100,00
Lifetime	
Physical Lifetime	10
Unit Costs	
Fixed Assets	
Capital Cost	Constant {1.000,00}
Calibration Period	Y1
Capital Cost Structure	
Switch off capex	1,00
Other Details	
Year Zero Data	
Max. Age of Installed Units	1
Installation Profile	Interpolated Series
	Period Y0 Q1
	Value 2,00 0,00
Normalise Profile	No



Neither capex nor depreciation for resources installed before the begin of model run

Solution: match depreciation calculation in STEM with accounting records (2)

- Use a separate resource to reflect depreciation of existing assets as in accounting records:
 - one 'asset resource' per resource or
 - one 'asset resource' for a group of resources (a platform or a separate part of the network)
- For that resource, a special depreciation rule is used which depreciates the resource if it is removed from the network (becomes redundant)
- With a decreasing number of deployment sites, one can remove the necessary number of resources from the network and (assuming, e.g., a capex = 1) generate the depreciation as it is stated by accounting records

Period	Y0	M1	Y1	Y2	Y3	Y4	Y5	Y6	Y7
Value	11.000,00	11.000,00	8.000,00	6.000,00	4.000,00	3.000,00	1.000,00	500,00	0,00

assets - Resource "Depr existing assets"/Other Details

Function: <none>

Label:

Pre-Run Installation

Max. Age of Installed Units: 1

Installation Profile: Interpolated Series { ... }

Normalise Profile: No

Planned Units: Interpolated Series { ... }

Redundant Unit Decomm. Prop.: Constant { 1,00 }

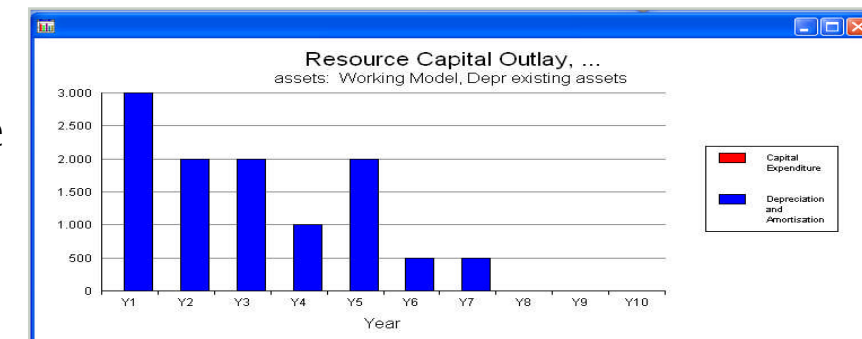
Churn Mode: Normal

Slack Equipment Bias

Use Slack: Even

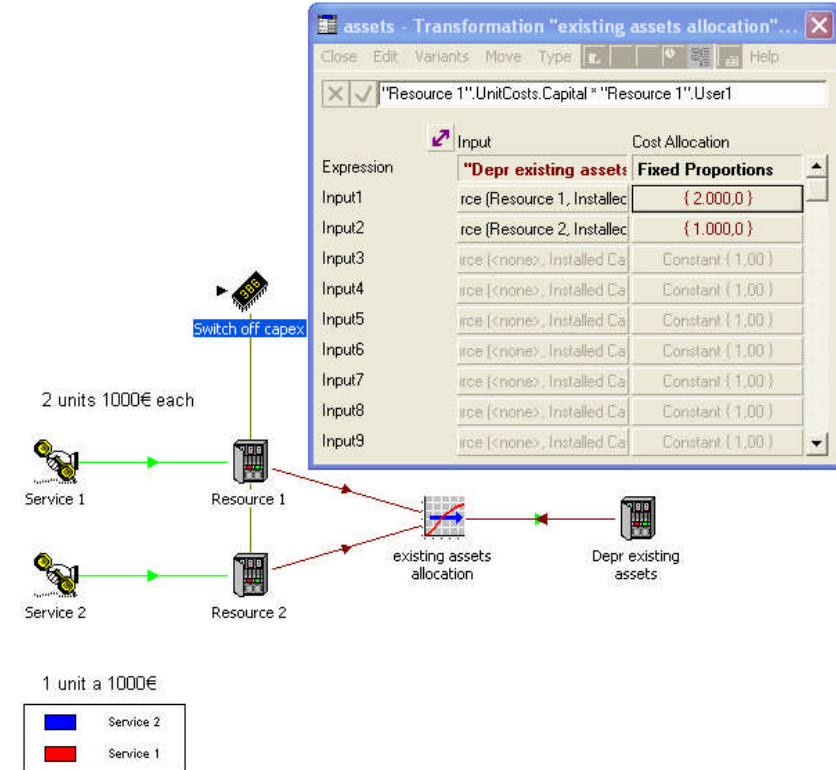
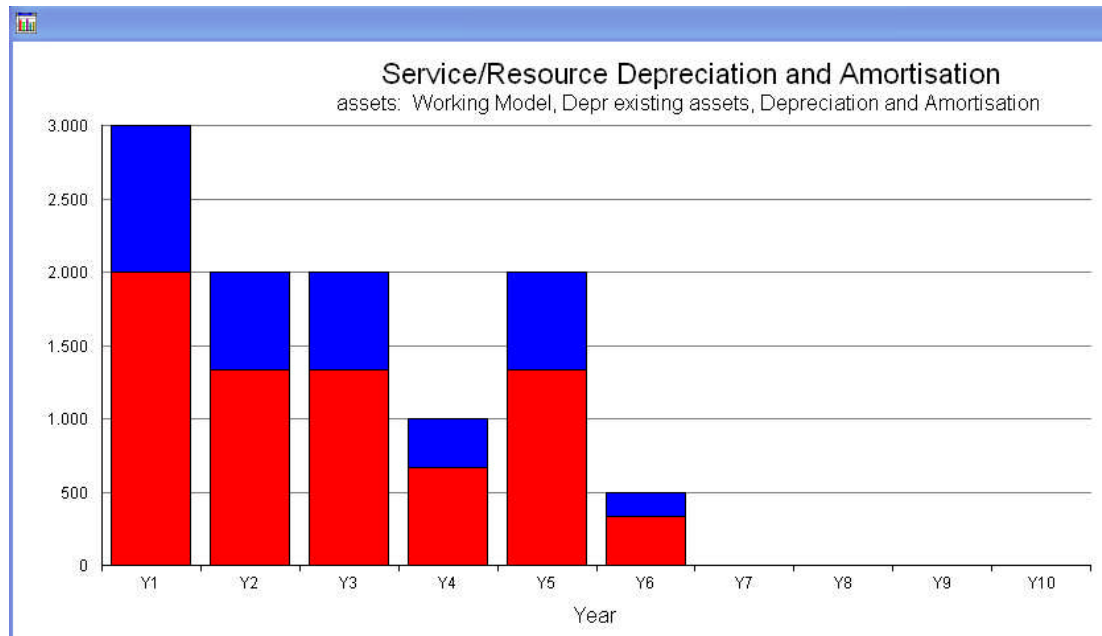
Make Slack: Even

Switched Off: No



Solution: match depreciation calculation in STEM with accounting records (3)

- Make sure that depreciation is processed in a way, that matches the purpose of the model, e.g.:
 - cost allocation to individual resources
 - apply utilisation



Limitations and suggestions: match depreciation calculation in STEM with accounting records

Limitations and disadvantages

- If the depreciation for a bunch of assets is reflected by one common asset resource, we are faced with a number of limitations which mainly arise from the unavailability of accurate Input data:
 - different age of assets is not considered (e.g., it can be that one of the resources is written-off within the first period, whereas the described approach will still allocate depreciation to it)
 - any replacement strategy would fail for the same reason
 - allocation to resources is based on a snapshot ratio

Advantages

- Simple and transparent
- Delivers results even if accurate (matching the model) records are not available
- Sufficient for costing and models with a short run time

Summary


- STEM's capex model and depreciation model is well-documented and readily understood and explained to colleagues and auditors
- It already benefits from 20 years of client request and refinement
- We constantly strive to keep the inputs and results relevant and up-to-date
- Your ongoing suggestions are therefore highly valuable!
- We do not guarantee to implement every last suggestion and detail

BUT

- We do aspire to be mindful of the underlying industry direction, to observe patterns in user requests, and to use our unique perspective to design and deliver innovative solutions which meet the essential challenges common to all

Platform evolution: priorities such as smarter results
browsing, driven by recent project experience

Robin Bailey – Managing Director, ‘Mr STEM’


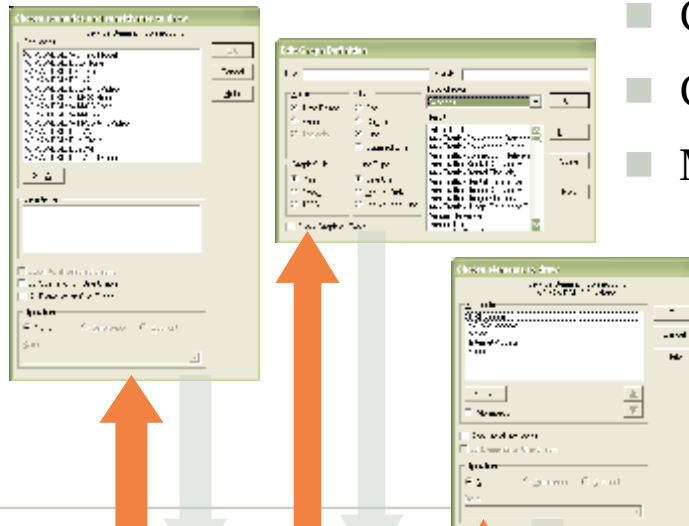


Introduction

- STEM development is, as ever, driven principally by user requirements, a virtuous combination of:
 - responding to specific requests and
 - designing more fundamental enhancements aligned with more general observed patterns of comment and requirement over time
- A couple of topical factors distinguish the current moment:
 - the formation of the new Implied Logic company puts product development back on the agenda with a renewed focus
 - a number of expert advisory projects in the last twelve months have yielded (through extended observation) a fresh perspective on:
 - what works and what doesn't for clients
 - how it feels to be a full-time user!
- Implied Logic has resolved to make a first new major release of STEM solely devoted to addressing the 'notes in the margin' from these recent projects

Theme I: smarter results browsing


- Wouldn't it be nice if you could:
- Double-click on labels to modify an existing graph?



- Change current or add extra scenarios
- Change current or add extra results
- Modify the selection of elements: *
 - with an automatic <all elements> option [facilitating a default workspace] for small models, or
 - show members of a collection(s)
 - sort by name or template variant
- Split 'complex' legends into two dimensions by colour and pattern:
 - for scenarios, elements or results
 - or elements × template variants
- 'Slice' individual charts or entire views via a drop-down (with option to lock)
- Filter individual charts or entire views by sub-string match on element names

What was that about graphing the members of a collection?

- We have long imagined accessing results from a cascading 'Results' icon menu-item *in the Editor*, especially:
 - Service / Connections
 - Transformation / Output
- So the same could work for a collection (or multiple selection) of elements



- And a modified chart could collapse to a new chart (or table) collection icon:*
 - drag further elements to add

- Results would only be displayed for the working model by default:
 - but a modified scenario selection would apply for new charts
- Showing a graph would automatically run the required scenarios:
 - <F5> could mean update visible
- Collections of variants or elements could be used to link variant or element selections across many charts:
 - somehow choosing between 'show members' and 'sum over'
 - the dimension icon means 'all scenarios' in a simple model
 - allow multi element-type charts
- The Editor would also be a much more friendly environment for defining new results by point-and-click reference

Integrating results views in the Editor

- There are many reasons why it would make sense to combine the Editor and Results applications:
 - ability to review inputs and outputs side-by-side
 - direct update on charts when elements are renamed or deleted
 - unified Excel input/output export
 - less application-window clutter!
- A new kind of view (or all Editor views) would need to accommodate charts 'in the view' to present multiple charts
- Editor views are tied to a single model, so this approach would only readily chart scenarios within the same model
- At present, the Editor has full 'process isolation' from any exceptions raised by the Results program
- Critically, the required actions to draw a graph must be *the same, or easier*, for this to be a step forward which will be welcomed by users
- So we should imagine how to draw :
 - *Connections* for one service
 - *Capacities and Installed and Incremental Units* for one resource
 - *Revenue* for all services
 - *Depreciation* for all resources
 - *NPV* for all scenarios
 - a new combination of results
- It should not be necessary to create any special icons to draw a graph:
 - in the first instance, a user should learn how this works by creating suitable collections from graphs

Theme II: just doing the right thing; phase 1

- Multiplier input for market segments and link with Connect tool
- Multiplier input for all transformations and show output unit in main dialog
- Transient, aggregate-mode resources and incremental transformations
- Limit for Monte Carlo deployment (e.g., carriers vs spectrum, 1-1 as special case) and path-protection multiplier
- Process icons in visible order rather than 'last clicked'
- Change 'unlink' action on market link to revert to zero to avoid <none> error
- Location Sites to go to button first, etc.
- Re-design time-series button metaphor
- Fix commands to view variant data from parameters (improve terminology in menu too, plus dialog-menu button)
- Suppress baffling 'Replace existing data' and 'Workspace merge' prompts!
- Load models from Windows Explorer
- Highlight errors before warnings, and add next/prev.-error toolbar buttons
- Suppress model name and some other details on chart minor titles
- Pre-select 'All elements on one graph' if/only if just one result and scenario
- Aggregation of quotient results, x3
- **Challenge: how many of these do you understand?**


Phases 2 and 3

- Replace *Unset* and *Unset Below*, *Freeze* and *Deep Freeze* with unambiguous *Unset* and *Freeze* commands
- Enable context menu (right-click) to copy/paste/freeze/unset dialog fields
- <F5> to repeat the most recent run command from the scenarios dialog
- Allow models to run in halves: H1/H2
- Allow one resource to drive another via an implicit transformation
- Find and replace for selection of icons and when copying a block of elements
- Animated icons
- Chart boundary tool
- Insert/move user data
- % operator
- Warn 5 days before soft-licence expires
- Look for STEMICON.DLL in c:\stem first, and create some fresh icons
- Zoom function in the Editor
- Option to create dialogs showing all and only the set inputs (like tooltips)
- Facility to create custom dialogs using point and click, cf. Create Graph
- Please comment on the suggested priorities

This space intentionally left blank

New market strategy: a 'small announcement to consulting partners' with a potentially major impact on future growth

Robin Bailey - Managing Director, 'Mr STEM'



Introduction

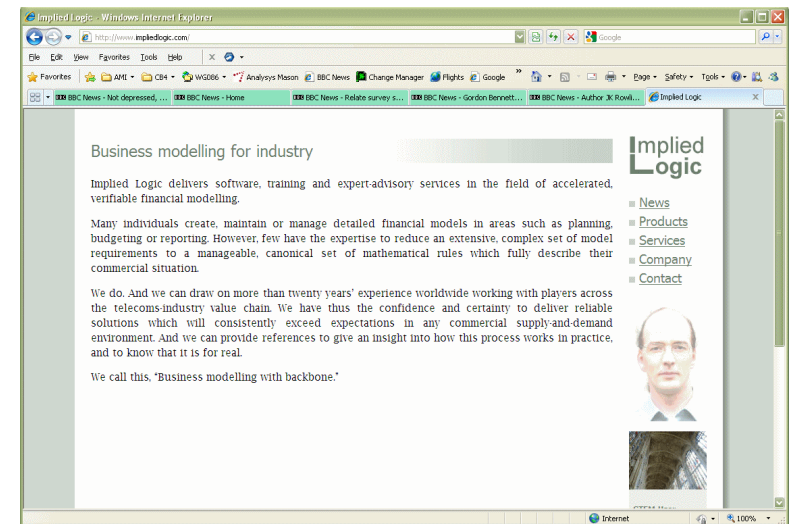
- Almost eight months ago, the wheels were already in motion to spin out the STEM business into a separate company from Analysys Mason (AM):
 - after abortive discussions at the start of 2009, agreement was reached early 2010 to have another look at the options
 - by mid-February, Robin Bailey was committed to exploring terms on which a suitable transaction could be effected and enlisted the moral support of a close group of associates familiar with the STEM business
 - after prolonged agonising over the name, *Implied Logic* was registered as a UK limited company on Friday 23 April
- So when this year's STEM User Group Meeting was first announced in the April newsletter, it was already an assumption that the event would become an Implied Logic event long before the guest list was finalised!
- When we alluded to a 'small announcement to consulting partners with a potentially major impact on future growth', what we meant was that *STEM would no longer be owned by a company with its own consulting interests*
- This session is an opportunity to understand our ambition in more detail

A change in name revisited

- From mid-February (and actually since before the previous STEM User Group Meeting in 2009), the working title for the new company was envisaged as *The STEM Expert Consulting Group*
- However, a concern emerged that this might be too long to be memorable
- As the future strategy crystallised and the clear benefit of being separate was better understood, the consulting word became a bit of a no-go*
- Literally scores of alternatives were considered before it struck Robin that:
 - “The core value of STEM is that users can quickly create and benefit from business models where the bulk of the calculation logic is implied from a clear and intuitive picture which they can create on-screen by themselves or in a workshop environment” (16 April)

Rapid progress from idea to reality

- The impliedlogic.co.uk / [.eu](http://impliedlogic.eu) / [.info](http://impliedlogic.info) / [.net](http://impliedlogic.net) / [.org](http://impliedlogic.org) and [.biz](http://impliedlogic.biz) Internet domains were registered on 18 April, followed a few days later by the alternative implied-logic.co.uk / [.eu](http://implied-logic.eu) / [.info](http://implied-logic.info) / [.net](http://implied-logic.net) / [.org](http://implied-logic.org) / [.biz](http://implied-logic.biz) and [.com](http://impliedlogic.com) domains
- *Implied Logic* was registered as a UK limited company on Friday 23 April, with an initial virtual office on the outskirts of Cambridge
- impliedlogic.com was acquired from a US incubator on Sunday 09 May
- An EU trademark application for *Implied Logic* was filed on Friday 14 May
- Implied Logic became the legal owner of the STEM business and associated intellectual property on Friday 18 June:
 - 21 years + 1 month from when Robin first started at Analysys!
- We have since registered both variants of [.asia](http://impliedlogic.asia) and [.co](http://impliedlogic.co)* domains



An uneasy mix of software and consulting

- Although originally a unique selling point of Analysys' consulting services in the early 90s, by 2000 STEM was already marginalised:
 - partly due to a decision long ago that AM consultants should *not* have incentives for developing STEM opportunities, compared to using Excel
- As clients increasingly demanded model-deliverables, rather than just reports, STEM became seen as a project 'problem', at least until D-STEM was ready in 2001, by which time the tide had turned
- STEM ceased to be core business, so there was no board-level focus, will to invest in development, or any significant visibility in company collateral
- After some challenging market conditions in 2008, I recognised a risk that STEM could be closed down over my head:
 - in spite of very positive client reactions I was continuing to experience
- As I started to develop a new wave of opportunities throughout FY2009/10, I decided it had become imperative to take a controlling interest in order to secure STEM's future stability and continued availability to the market



Reasons to leave ...

- For some time I had wondered if the AM brand sat comfortably with STEM:
 - my clients have typically had very specific intentions for STEM
 - and have not wished to pay a premium for AM consultants to advise on their strategy (as opposed to how to achieve their modelling aims)
- More fundamentally, I began to realise that many external *consultants* could benefit enormously from the credibility, reliability and large-model scale which STEM offers:
 - were it not for their evident concerns about sharing the specifics of commercial opportunities with a competing consulting company (a significant historical conflict of interest)!
- In fact we had worked with a small number of independent consultants who would licence STEM on a project basis from time to time:
 - but all of these opportunities were based on previously established trust with Robin (former clients or colleagues)
 - and this was never going to turn into a growth business while the general conflict as identified above persisted

... and reasons to thrive independently

- Our brand focus will be on the STEM platform(s) as an enabling technology, and the promotion of expert practitioners who can help clients achieve excellence with our software (rather than offering strategic solutions)*
- Our mission is to gain recognition, credibility and more significant market adoption of the STEM modelling process through collaboration with a network of consulting partners worldwide
- By promoting geographical and industry diversity, our aim is for such partners to help introduce the modelling paradigm into other industry sectors*, and in so doing, to become our principal sales channel (so that we can focus on what we know best)
- Our principal revenues will come from software, and our complementary services will be positioned around training and expert advice:
 - never seeking to engage in the kind of end-to-end consulting projects which might alienate or deter our intended consulting partners
 - our proposition is mathematical clarity and modelling certainty

The last six months have not gone exactly as anticipated

- My expectation was that it would take some planning and support to effect a commercial transition out of AML
- So my original intention was to spend the first few months cementing a suitable partnership from within my network of external associates ...
- ... before then engaging in a potentially prolonged negotiation with AML
- One thing I was clear about was the importance of being able to announce the new entity at the STEM User Group Meeting in September 2010:
 - or early October as it turned out to be when I found that my preferred dates were already taken at King's!
- In practice, it has been rather different:
 - first, it turned out to be *much* harder than expected to figure out terms to be in business with such long-established friends and arms' length consulting partners, and
 - second, my initiative just so happened to coincide with a drive to streamline AML operations and simplify its marketing proposition
- In spite of STEM being very profitable in 2009/10, AML was ready (and you could say even impatient) to negotiate
- So by Friday 18 June, Implied Logic was the proud owner of STEM:
 - and now we are gradually building up the support organisation

Implied Logic now and in the future

- After several project interludes and a summer holiday, on Monday 16 August Implied Logic moved into its serviced office co-located with the initial, virtual office address on the Vision Park in Histon on the outskirts of Cambridge
- Our initial website and email was hosted c/o our friendly home ISP, but by mid-September, these services had been migrated to our new, virtual IT assets (in a cloud some place)
- In and around a number of ongoing projects, we have taken on:
 - *Beth Stevens*, part-time Production Assistant in Cambridge, UK
 - *Mike Waterman*, part-time Support Engineer in Wellington, NZ
- A network of [currently] external associates may provide services on our behalf under sub-contract, including:
 - *Frank Haupt*, Expert Practitioner in Berlin, who is representing us during this User Group Meeting
- We have also out-sourced accounts and HR support from conveniently located and trusted friends and neighbours
- We plan to recruit a small software team as fast as we dare and is practical
- We will work closely with clients in the next three years, learning from current failings and gradually removing perceived 'roadblocks' to learning
- Incremental enhancements to STEM will continue up to the planned launch of a *New STEM* product range (desktop and online) in September 2015

Legal and technical implications

- As the new owner of the IPR in STEM, Implied Logic is the only legal entity which can effect a software licence ...
- ... so Implied Logic has automatically assumed the licensor role in all existing licence agreements
- However, Analysys Mason retains certain warranty and maintenance obligations which will only cease when the relevant contracts are renewed in due course with Implied Logic:
 - so Implied Logic has accepted a back-to-back obligation with AML to perform all the relevant support duties on AML's behalf
- Fortunately the support team has not changed, so you are in good hands, via support@impliedlogic.com
- STEM 7.2b is hard-wired to request soft licences from www.analysysmason.com
- The same server functionality is now ready on www.impliedlogic.com, and we will soon release a STEM 7.2c which talks to the new server (and no longer mentions AM!)
- After a short notice period, we will decommission the AM soft-licence server, and then complete the removal of all remaining STEM property from the AM office at St Giles Court in Cambridge




Company style

- The majority of the audience at the STEM User Group Meeting are rather familiar with Robin Bailey and will recognise his values in Implied Logic:
 - openness, trust, integrity and respect
 - smart, original solutions
 - small pieces of joined-up logic providing a rock-solid foundation for client business-modelling endeavours
 - zero-tolerance attitude towards bugs and dumb processes
 - and rather fond of the colour green – since long before environmental responsibility became fashionable
- Any questions – please ask
- Thank you for your attention today, and for widespread support and encouragement over the roller-coaster changes of the last six months
- We hope that you are going to like working with Implied Logic!

This space intentionally left blank

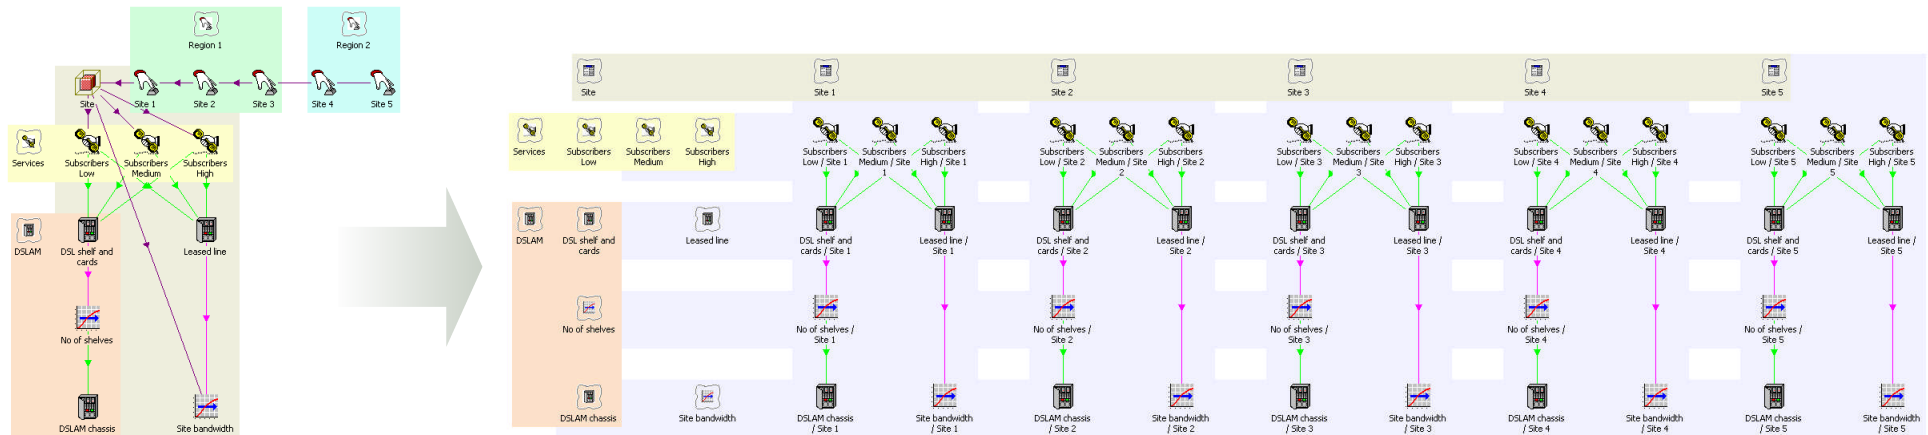
Showcase: a jaw-dropping technique, and some personal recommendations on linking style

Robin Bailey – Managing Director, ‘Mr STEM’



A quick introduction to templates in STEM

- Many of you will be familiar with the concept of *template replication*:
 - in essence, the ability to systematise the process of creating parallel, named copies a chunk of model structure with varying assumptions
 - without having to do it by hand
 - with the compelling property that any subsequent changes to the template structure are automatically propagated across all copies

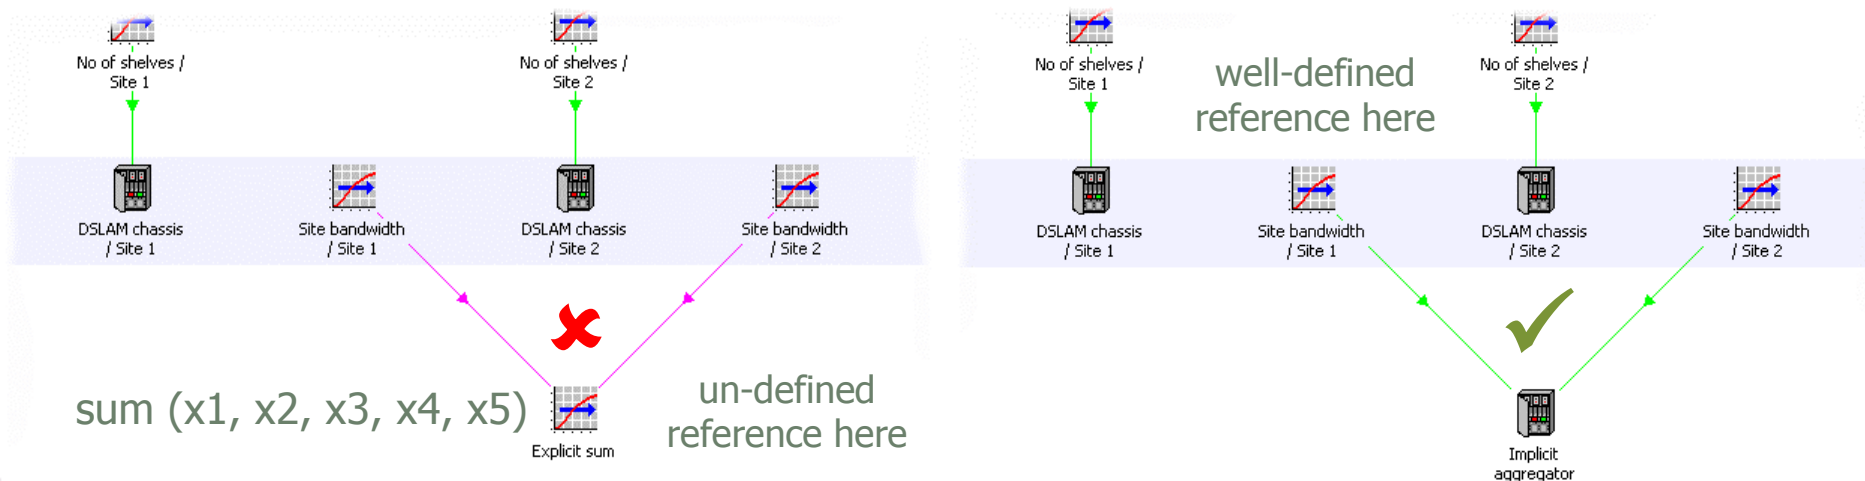


template + n variants

n consistent copies generated at run-time


Aggregating output from template variants

- A few of you may recall from training, or from your own projects, the slightly unobvious question of how to drive demand from a template onto a core, shared network infrastructure, or any common system in general
- The issue is that the obvious summing transformation cannot look back at the individual variants – because the template copies are not instantiated until run-time
- So the training courses introduce the technique of driving demand from a template output element onto a common (non-replicated), ‘aggregator’ resource, which supports the required many-to-one relationship



Pushing template replication one step further

- So imagine the challenge when a client on a recent project asked:
 - *“Is it possible for a template to ‘drive itself’?!?”*
- This sounds circular in the extreme until you know more about the context:
 - but the actual model is commercially sensitive and client confidential
 - so instead we will look at a simple and contrived example which illustrates the same technique, even though this particular problem could easily be solved by a less sophisticated method!
- Consider a short, linear network, which uses OADMs* to add and drop optical paths at each of the nodes along a trunk route



* optical add/drop multiplexers

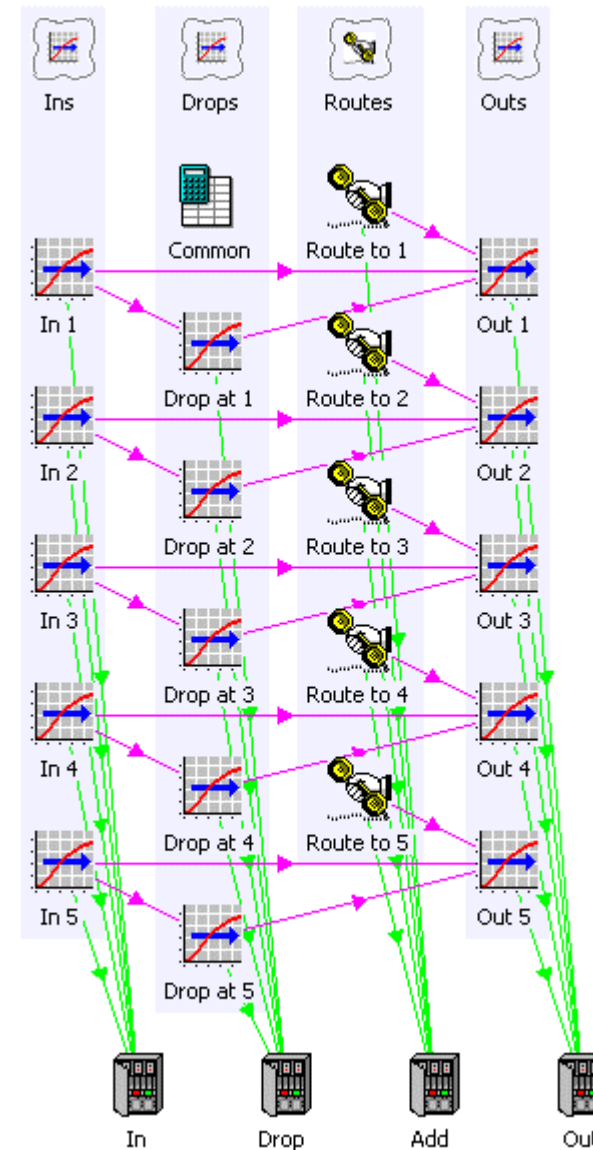
Wiring variants in series rather than in parallel

- From a set of circuit requirements for each logical route (pair of end points), we seek to dimension virtual add, drop and through capacity at each node
- This might translate to physical ports, but the specifics are not critical
- What we would like to do is structure a model such that we can:
 - feed in at one end a traffic matrix like the one shown opposite
 - have a node template calculate the relevant adds and drops and residual through traffic
 - pass this from one node to the next to dimension all the nodes
 - without the model being circular!

Circuits	A	B	C	D	E
A					
B	42				
C	11	17			
D	35	21	16		
E	28	19	32	25	

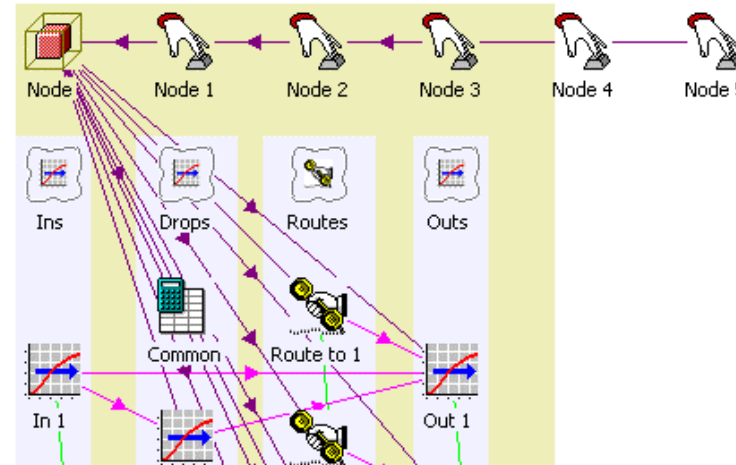
Calculating add, drop and through capacity

- We need to build a generic template of a node which may require circuits for any end-point:
 - even if a given node will only count circuits to those nodes following it in the table
- Assume that we have an incoming requirement in *Circuits* destined for each of nodes 1-5 (i.e., A-E)
 - working from A-E, we will never actually count any circuits terminating at A
- All incoming circuits require inbound interfaces
- Only those terminating at a given node require drop interfaces, so we have to check against the node number, stored as a template parameter
- All routes originating at a given node require add interfaces, driven by separate services for each end-point with demand varying by node
- Outbound interfaces are required for the sum of added circuits and through circuits, the latter calculated as Inbound Circuits less Drop Circuits



Template replicated for each of the five nodes

- This whole calculation structure is wrapped up as a template:
 - with parameters identifying the node number and originating circuit numbers
 - and variant elements naming each of the five nodes and providing specific values for each parameter




ROADM v04 - Template "Node"/Variant Data

Close Edit Variants Move [Icons] Help

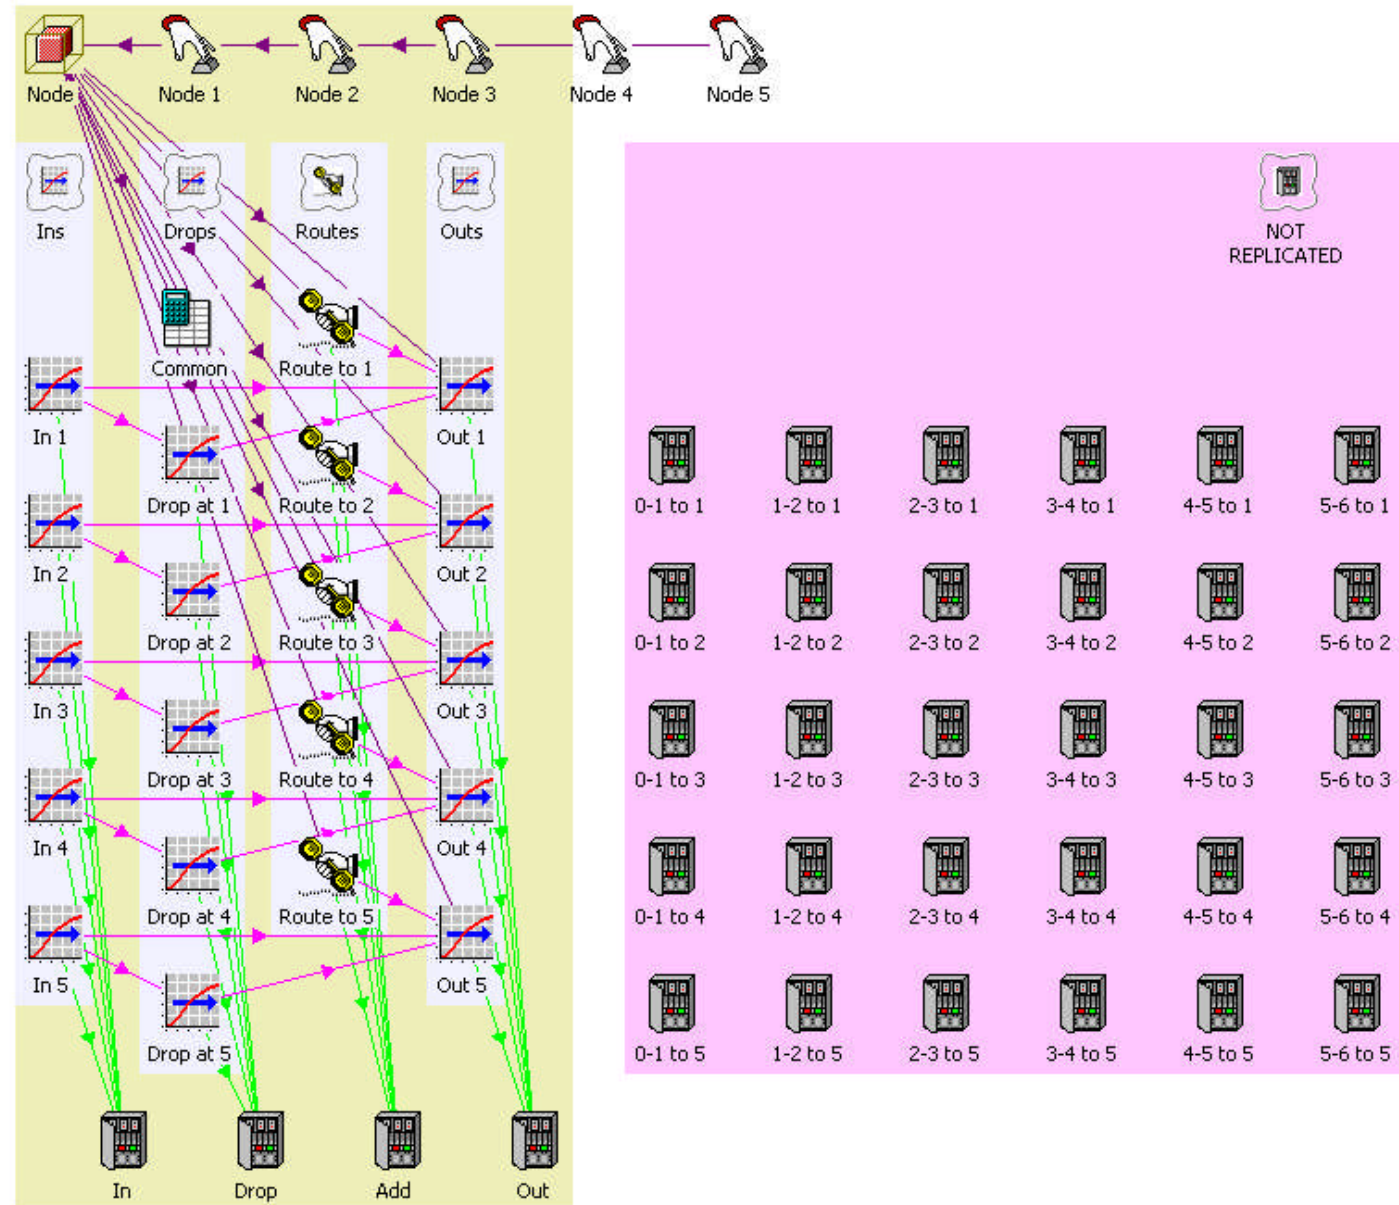
	Node 1	Node 2	Node 3	Node 4	Node 5
Routeto1.CustomerBase	Constant { 0.00 }	Constant { 0.00 }	Constant { 0.00 }	Constant { 0.00 }	Constant { 0.00 }
Routeto2.CustomerBase	Constant { 42.00 }	Constant { 0.00 }	Constant { 0.00 }	Constant { 0.00 }	Constant { 0.00 }
Routeto3.CustomerBase	Constant { 11.00 }	Constant { 17.00 }	Constant { 0.00 }	Constant { 0.00 }	Constant { 0.00 }
Routeto4.CustomerBase	Constant { 35.00 }	Constant { 21.00 }	Constant { 16.00 }	Constant { 0.00 }	Constant { 0.00 }
Routeto5.CustomerBase	Constant { 28.00 }	Constant { 19.00 }	Constant { 32.00 }	Constant { 25.00 }	Constant { 0.00 }

- But how is one instance of this template wired up to the next?

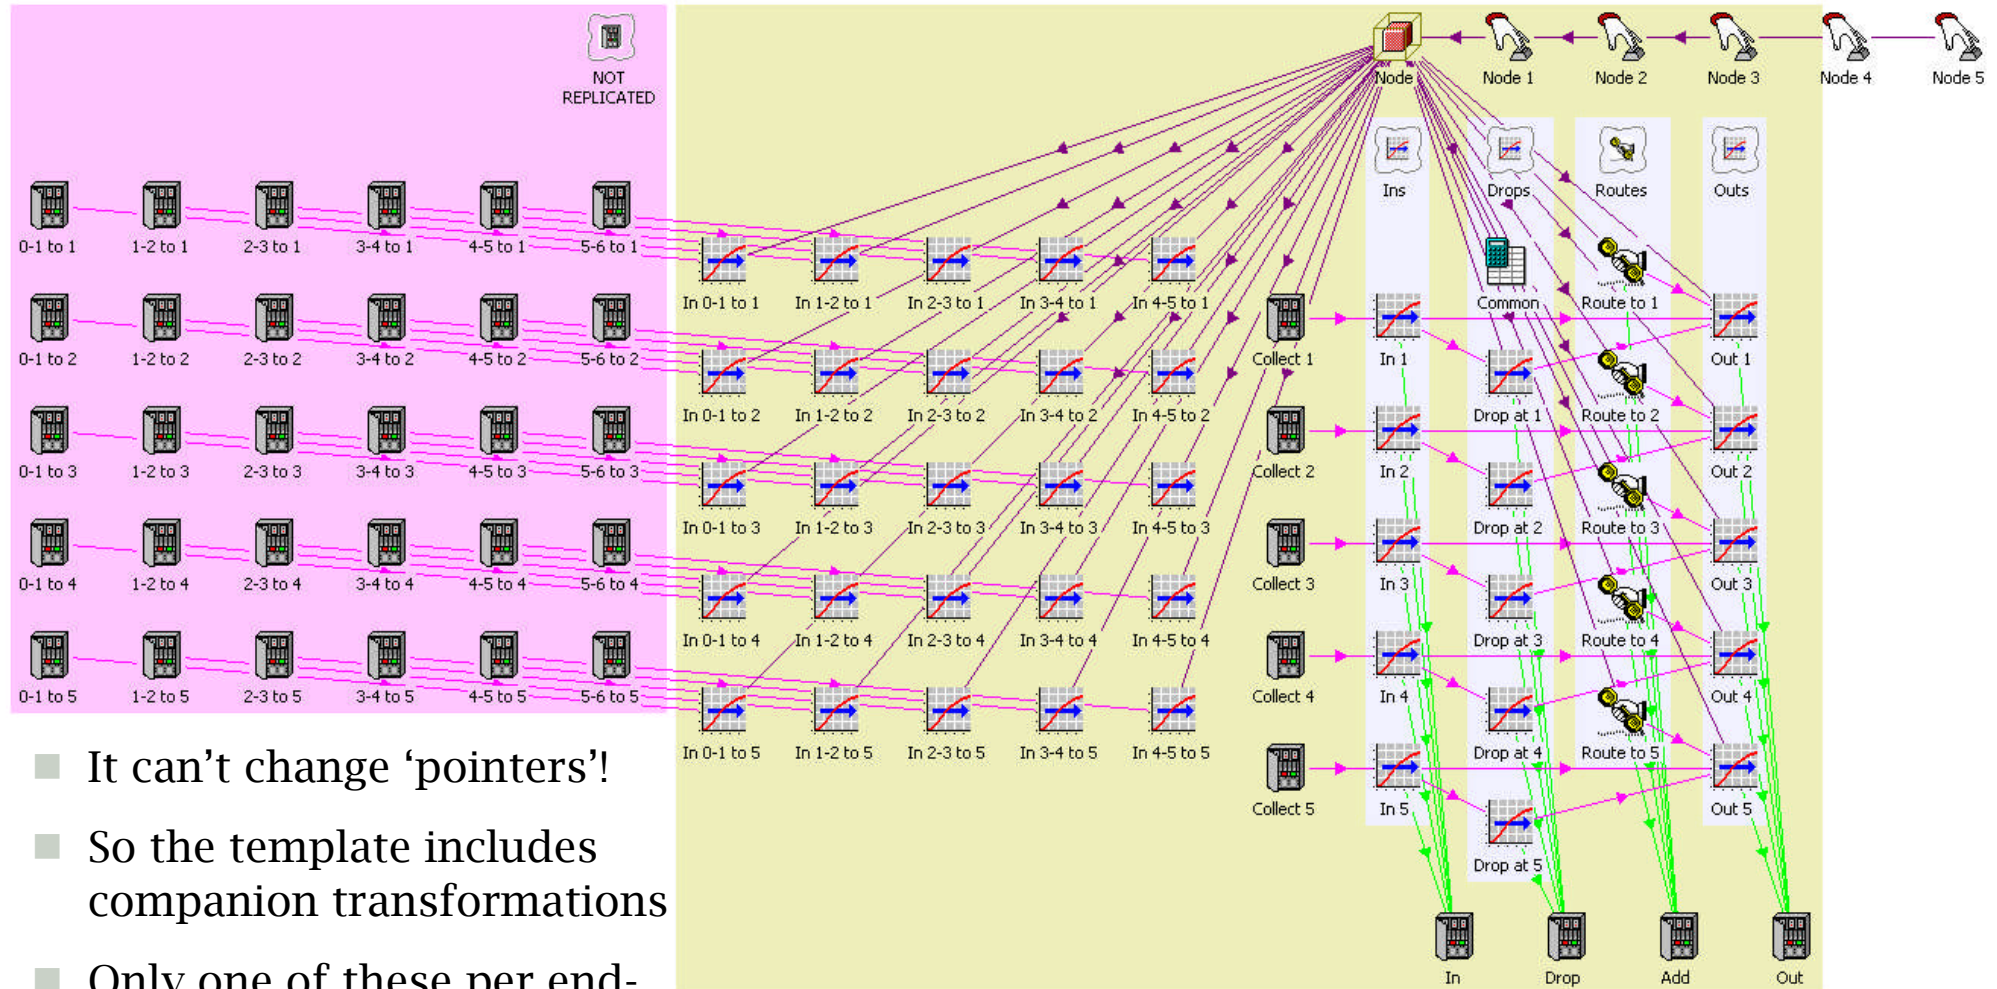


Outbound circuits to intermediate totals

- Recall the earlier example of how to aggregate the output of a template
- It is not hard to guess that you could drive demand out of the template onto a non-replicated resource which could in turn feed another instance
- The template could map to different resources by node (effectively the links)
- But how can the template dynamically select the correct link resource as input?



Intermediate totals to inbound circuits



- It can't change 'pointers'!
- So the template includes companion transformations
- Only one of these per end-point is mapped to a 'collector' resource which drives the inbound demand for that end-point on a given node

Add, drop and through capacity by node

STEM Results | Implied Logic - [ROADM v04;DEFAULT;ROADM v04.smr]

File Edit View Graphs Format Consolidation Config Options STEM Help

Service Demand - Connections (ROADM v04: Working Model, Connections)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Route to 5 / Node 1	28	28	28	28	28	28	28	28	28	28	28
Route to 4 / Node 1	35	35	35	35	35	35	35	35	35	35	35
Route to 3 / Node 1	11	11	11	11	11	11	11	11	11	11	11
Route to 2 / Node 1	42	42	42	42	42	42	42	42	42	42	42
Route to 5 / Node 2	19	19	19	19	19	19	19	19	19	19	19
Route to 4 / Node 2	21	21	21	21	21	21	21	21	21	21	21
Route to 3 / Node 2	17	17	17	17	17	17	17	17	17	17	17
Route to 5 / Node 3	32	32	32	32	32	32	32	32	32	32	32
Route to 4 / Node 3	16	16	16	16	16	16	16	16	16	16	16
Route to 5 / Node 4	25	25	25	25	25	25	25	25	25	25	25

Resource Used Capacity (ROADM v04: Working Model, Used Capacity)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Out / Node 1	116	116	116	116	116	116	116	116	116	116	116
In / Node 1	0	0	0	0	0	0	0	0	0	0	0
Add / Node 1	116	116	116	116	116	116	116	116	116	116	116
Drop / Node 1	0	0	0	0	0	0	0	0	0	0	0

Service Demand - Connections (ROADM v04: Working Model, Connections)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Route to 5	104	104	104	104	104	104	104	104	104	104	104
Route to 4	72	72	72	72	72	72	72	72	72	72	72
Route to 3	28	28	28	28	28	28	28	28	28	28	28
Route to 2	42	42	42	42	42	42	42	42	42	42	42

Resource Used Capacity (ROADM v04: Working Model, Used Capacity)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Out / Node 2	131	131	131	131	131	131	131	131	131	131	131
In / Node 2	116	116	116	116	116	116	116	116	116	116	116
Add / Node 2	57	57	57	57	57	57	57	57	57	57	57
Drop / Node 2	42	42	42	42	42	42	42	42	42	42	42

Service Demand - Connections (ROADM v04: Working Model, Connections)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Routes / Node 1	116	116	116	116	116	116	116	116	116	116	116
Routes / Node 2	57	57	57	57	57	57	57	57	57	57	57
Routes / Node 3	48	48	48	48	48	48	48	48	48	48	48
Routes / Node 4	25	25	25	25	25	25	25	25	25	25	25

Resource Used Capacity (ROADM v04: Working Model, Used Capacity)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Out / Node 3	151	151	151	151	151	151	151	151	151	151	151
In / Node 3	131	131	131	131	131	131	131	131	131	131	131
Add / Node 3	48	48	48	48	48	48	48	48	48	48	48
Drop / Node 3	28	28	28	28	28	28	28	28	28	28	28

Service Demand - Connections (ROADM v04: Working Model, Connections)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Routes / Node 1	116	116	116	116	116	116	116	116	116	116	116
Routes / Node 2	57	57	57	57	57	57	57	57	57	57	57
Routes / Node 3	48	48	48	48	48	48	48	48	48	48	48
Routes / Node 4	25	25	25	25	25	25	25	25	25	25	25

Resource Used Capacity (ROADM v04: Working Model, Used Capacity)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Out / Node 4	104	104	104	104	104	104	104	104	104	104	104
In / Node 4	151	151	151	151	151	151	151	151	151	151	151
Add / Node 4	25	25	25	25	25	25	25	25	25	25	25
Drop / Node 4	72	72	72	72	72	72	72	72	72	72	72

Service Demand - Connections (ROADM v04: Working Model, Connections)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Routes / Node 1	116	116	116	116	116	116	116	116	116	116	116
Routes / Node 2	57	57	57	57	57	57	57	57	57	57	57
Routes / Node 3	48	48	48	48	48	48	48	48	48	48	48
Routes / Node 4	25	25	25	25	25	25	25	25	25	25	25

Resource Used Capacity (ROADM v04: Working Model, Used Capacity)

Circuits	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
Out / Node 5	0	0	0	0	0	0	0	0	0	0	0
In / Node 5	104	104	104	104	104	104	104	104	104	104	104
Add / Node 5	0	0	0	0	0	0	0	0	0	0	0
Drop / Node 5	104	104	104	104	104	104	104	104	104	104	104

Observations and a planned simplification


- This technique is not for the beginner!
- And it is an absurdly complicated way to calculate results for this specific, contrived example, as mentioned at the beginning
- One of the tricky aspects is that:
 - not only do you have to avoid the model being circular
 - dummy drivers and sinks have to be provided for all the elements which are only linked at one node
 - one of these dummies has to be replicated too
- Of course the results are much richer if the requirements vary with time
- Currently, a mapping is regarded as significant if it is defined by a formula:
 - so each switched mapping [50] must be an explicit template parameter!
- It would be *much* more economical if a zero formula were disregarded
 - every mapping could then be defined as a function of the node number

Personal recommendations on linking style

- If an assumption is common to several elements, it is often helpful to use a location element for 'module data':
 - with a different icon of choice
 - makes it easy to find assumptions
 - and convenient to link from Excel
 - still readily copied with associated elements (unlike global user data)
 - use a Cost Index for beginning-aligned inputs such as tariffs
- Generally better to link from module user data to per element user data:
 - so formula references are local
- Nice example is a Location called XR with fields such as GBP, EUR and USD
 - then you can enter `XR.GBP * 42`
- If a transformation converts from channels to bandwidth, don't call it 'Voice bandwidth per channel':
 - this is not what it calculates (and what appears on results charts)
 - but how it is calculated
- Better to reflect the absolute output:
 - e.g. 'Voice bandwidth'
 - link the multiplier to User1
 - named as 'bandwidth per channel'
 - very clear in tooltip summary
- Both of these practices were honed on client projects in the last 12 months
- Several clients were keen to make elements of the models easy to re-use

Opex work-out: consistent planning of maintenance, support staff and overhead costs

Robin Bailey – Managing Director, ‘Mr STEM’



Introduction

- Capex and depreciation are the usual mainstay of BoE* cost models
 - it is one thing to map required demand to equipment capacity
 - but quite another to estimate or model the associated human resources
 - even though these may account for more than 50% of the total cost!
- We all adopt the usual 'guesstimate' fudges if time or data is limited:
 - total opex as a % of revenue, or
 - bottom-up opex as a % of capital cost per line item
- Far better is to model the processes around installation and maintenance:
 - plan, build, operate, test, repair
 - staff, tools, spares, vehicles, buildings
- But there is a hidden gotcha waiting for when you least need it:
 - how to interpret activity flows in a quarterly or monthly context?

Working with a 'guesstimate' fudge

- Benchmarks always provide useful magnitude checks:
 - e.g.; calculate opex as a % of revenue for a 'similar' business
 - then check that the result for your model is between $x/2$ and $2x$
 - if not, then look for concrete reasons to explain the difference
- But if you have no time or other rules from which to calculate an opex line ...
- ... then in an emergency you can use the benchmark % to estimate a result:
 - $myOpex = myRevenue \cdot x$
- This is nothing better than admitting that your model has an unknown domain of scale x compared to revenue
- Slightly better is to break this down by asset type ...
- ... and to say that, for a given group of assets, i , there is a specific benchmark, y_i , which indicates opex such as maintenance cost as a % of capex
- Then you can build up an estimated opex result summed over asset types:
 - $opex = \sum_i capex_i \cdot y_i$
- This has the benefit of focusing opex sensitivity around assets for which opex is generally more significant ...
- ... but assumes opex is proportional to installed base (e.g., no learning curve) ...
- ... and still admits an unknown overhead, albeit unknown per asset type

Modelling the processes around installation and maintenance

ACTIVITIES

- Opex mostly arises from various activities in an asset lifecycle:
 - plan
 - build
 - operate
 - test
 - repair
- This has a near identical structure to the more familiar model for costs driven by service requirements
- We just have to trade the notions of instantaneous connections and capacity demands for *task rates* * required to fulfil a given activity



RESOURCES


- The cost relates to the resources required to fulfil these activities:
 - staff
 - tools
 - spares
 - vehicles
 - buildings
- An opex resource may have a capacity defined as an effective task rates:
 - e.g., trouble tickets per month
- The unit cost is directly the 'time cost', which is proportional to the cost per activity via the task rate
- This is best understood by example!

Using task estimates and resources to forecast headcount

- Consider a Switch resource which requires regular maintenance to attend to recurring faults or typical unplanned incidents
- This is a function of the the total installed base, and relatively easy to model, using two basic ingredients
- A transformation uses assumptions for:
 - the number of faults per annum
 - the average number of hours required to correct a fault
- ... to calculate an annual fault hours result (i.e., instantaneous annual rate)
- The capacity of an Engineer resource is derived (as a matching instantaneous annual rate) from assumptions for:
 - weeks per year
 - hours per week
 - effective utilisation
- This works fine, even in a quarterly context, because only a quarter of the salary cost is charged each quarter
- An increasing salary can be modelled with a cost trend; and a learning curve effect could be modelled within the Average Correction Hours per Fault input, factoring in a floor and (decreasing) multiplier
- The engineer may also be shared by different types of equipment.



Annual task hours and required engineer capacity


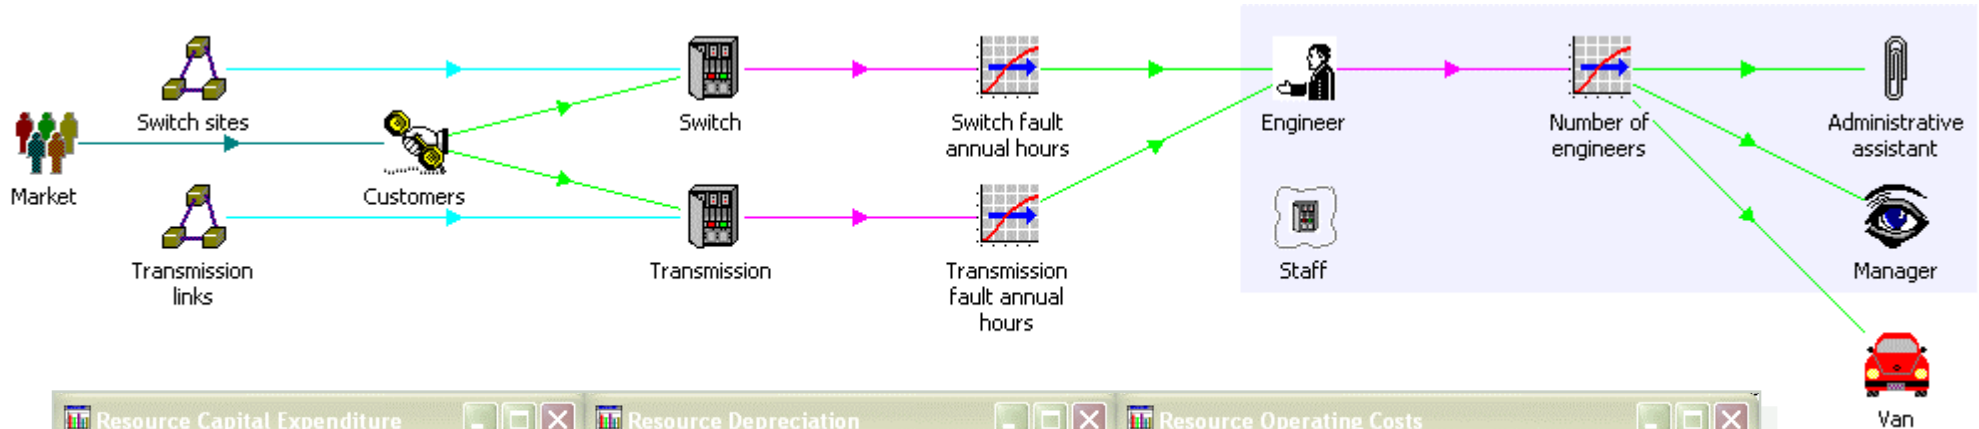


- An engineer working 48 x 37.5 hours a year, at a utilisation of 80%, could result in some of the 1440 annual hours being left slack (i.e., unused)
- Or if the engineer is shared from a company pool, then you could use a lower hours-per-week assumption to define a lower capacity
- The cost for a sub-contracted engineer could be defined as an *operations* cost, only arising from actual usage:
 - but the hourly rate might be higher for this 'pay as you go' approach
- Scenarios may be used to evaluate best and worst-case fault-rate assumptions

Adding dependent staff and support assets

- The standard technique of using a transformation with a basis of resource installed units (engineers) can be used to model a variety of dependencies from the engineer:
 - administrative support
 - manager
 - vehicle
- Each of these resources may be characterised with a capacity measured in (supported) engineers, e.g.:
 - one admin per five engineers
 - one manager per ten engineers
 - one van per [one] engineer
- In a more complicated model, these secondary support resources might be shared by different types of engineer
- Or the engineer resource might be shared by different types of equipment
- Some of these resources may need to be separately based at more than one physical location, readily modelled with the usual STEM deployment concept:
 - engineer resource required at multiple support sites or centres
 - admin resource also per site
 - manager resource may be centralised (i.e., shared by sites)

Secondary support resources and opex breakdown

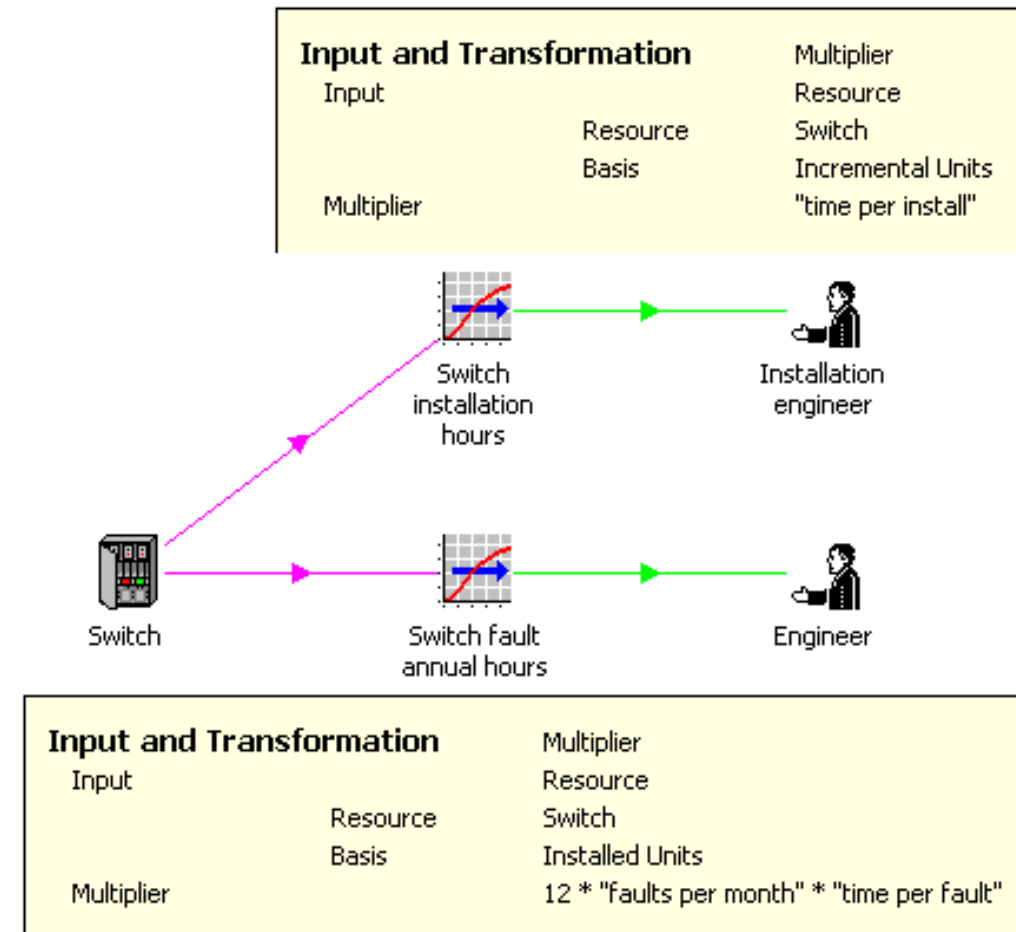


Handling short-term contracts

- One issue, which arises if demand is intermittent or falling, is that a STEM resource has a specific lifetime:
 - currently at least one year
- So you may end up with an engineering cost [resource] for a few quarters after a switch is no longer required:
 - the same issue would apply to customer support staff resources if demand is falling each quarter
- One potential fix for this would be to allow a lifetime less than one year:
 - but it is not clear how this would behave in a model running in years
 - install more than once?
- It might be preferable to class a resource as a *transient asset* which would disappear as soon as it is not required (perhaps equivalent to setting Redundant Unit Decomm. Prop. = 1):
 - or, better, to have a release time which could model a notice-period
- Another approach would be to think of the capacity of a resource in terms of:
 - an absolute *number of tasks that can be done for a certain cost*, and
 - for such a resource to persist for the current period only
- This idea is explored in more detail later in the more general context of incremental effects

Quantifying installation-engineer time in an annual model*

- So far we have examined activities related to total demand, matched to an instantaneous annual effort rate
- For models running only in years, it is straightforward to model equipment installation rates and required manpower levels too, again in two stages:
 - incremental units of the Switch resource can be scaled by hours per installation to calculate the annual hours required
 - an Engineer resource derives its capacity (as a matching instantaneous annual rate) from user data assumptions for weeks per year, hours per week and effective utilisation (just as above)

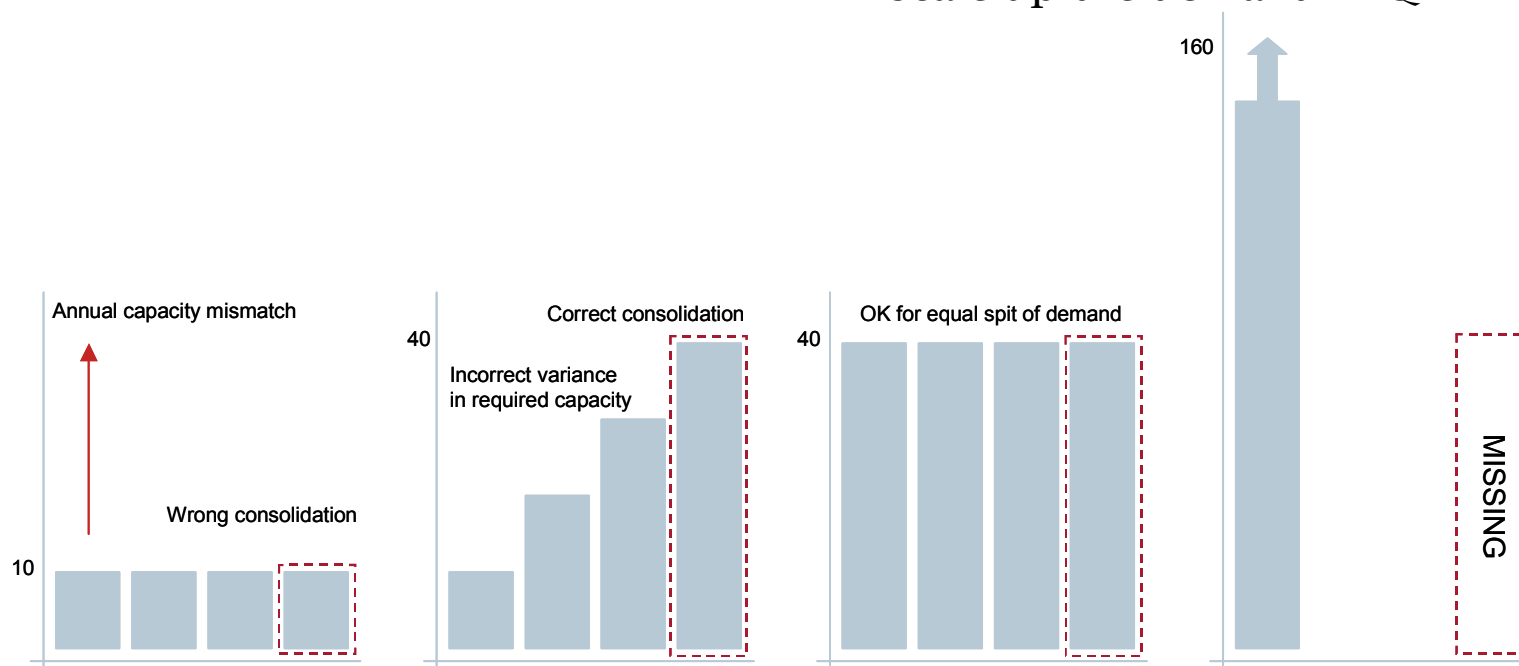


Interpreting incremental quantities in quarters or months

- Consider two cases:
 - one unit installed in each quarter
 - four units all installed in Q1
- In the first case, if we want to use a resource with a capacity specified as an annual rate, we need to scale up by factor of four in a quarter to get the same share of a permanent resource
- I.e., we must demand 40 hours in a year if we are to have 10 hours in each quarter:
 - bear in mind that only a quarter of the salary is charged each quarter
- This can be achieved by using the `periodLen()` function in a formula for the multiplier
- In a case where all the incremental units are in the first quarter, then this would be fine if the resource capacity did not persist beyond the quarter
- It makes sense that 40 hours would be required in that quarter: potentially more staff are being employed on a shorter contract
 - some additional sketching may be required to explain this!
- But how should an incremental units transformation Output result be consolidated (i.e., presented in annual format when a model has been run in quarters or months)?

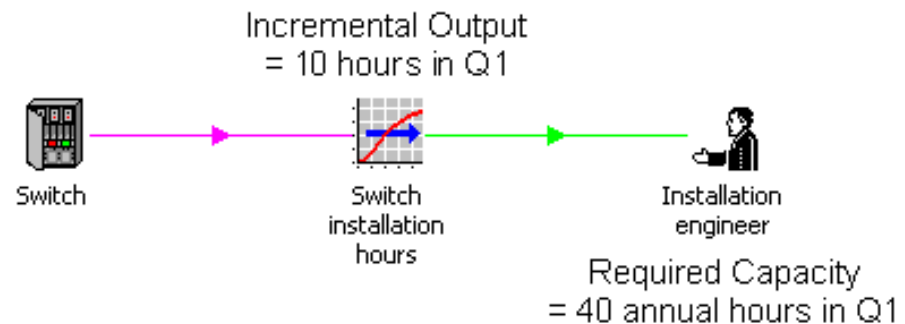
Alternative consolidations of quarterly incremental demand

- In general, the Output of a transformation is assumed to be instantaneous (as per most demands in STEM), and so we take the value from the final period in a year:
 - but this would completely hide the transient in the case of everything being installed in Q1
- So STEM is currently coded to accumulate an incremental units basis during a year specifically to avoid this problem:
 - but then this leads to too much capacity of the resource being required in later quarters, if the periodLen() approach is used to scale up the demand in Q1



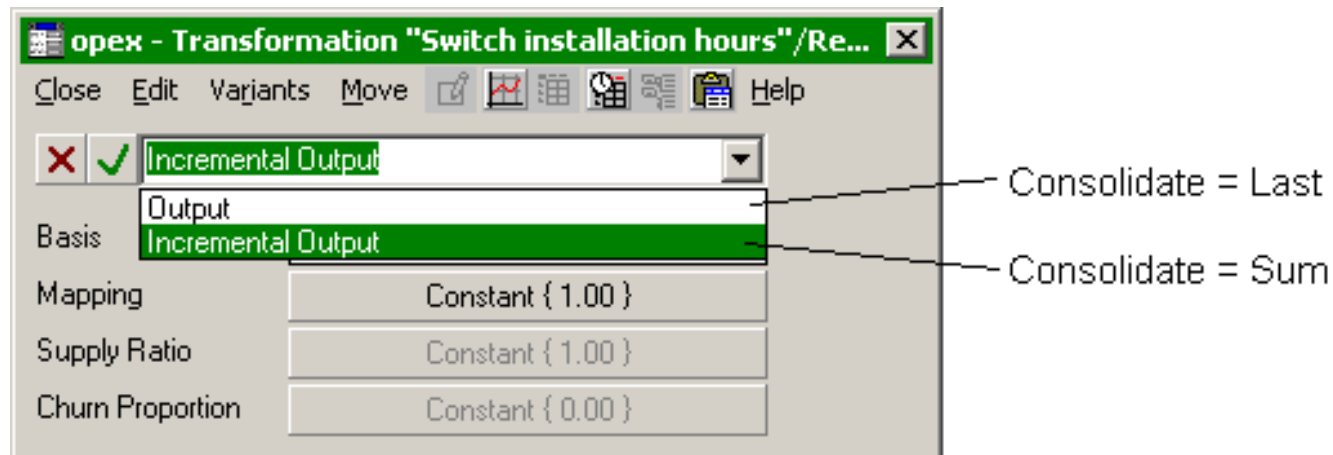
A potential solution for comment

- Suppose instead that we were to introduce an aggregate-mode transformation (or output), which would not cumulate over the year, and which would be consolidated by summation
- This would not be able to use the periodLen() approach, as the output for a quarter should relate only to the increment in that quarter (and would be summed across quarters on consolidation) ...
- ... but STEM could understand that the periodLen() scaling should be *built in* when matching such an aggregate demand to instantaneous capacity of a resource
- That is, the demand (derived from incremental units) would be 10 hours in the quarter, and this would turn into a requirement for 40 hours per year instantaneous capacity of a resource
- Or we could also introduce an aggregate-mode resource whose costs would be associated with aggregate demand in the current period only, and not scaled as a proportion of a year
- This is the concept of a 'transient' resource mentioned earlier



Transformation output bases and related consolidation modes

- For the Results program to know how to consolidate such an aggregate-mode output, we could have two types of transformation
- More manageably (and with much more general clarity), we could add a second, Incremental Output basis for transformation outputs:
 - or we might call the alternatives *Aggregate* and *Instantaneous*
- So you would choose which basis was being used when linking from a transformation




Would it ever make sense for both modes to be available?

- For single input transformations, you could infer the mode from its input
- Whereas an expression transformation would have to specify the user's intention
- The resource Incremental Units basis would be an aggregate input, as would:
 - service traffic (volume)
 - service revenue
- Note: both the latter are currently scaled as Annual Traffic and Annual Revenue bases:
 - they reflect the instantaneous traffic/revenue per year
 - as opposed to actual traffic or revenue in the current period
- The 'other' mode would be left at zero, or perhaps more usefully:
 - a standard output could be cumulated from incremental output within a year as per the current incremental units basis
 - an aggregate output could be inferred as a delta of two consecutive standard outputs
- This needs more thought!
- However, recognising these two bases and classifying them properly would enable STEM to make more meaningful and reliable incremental calculations ...
- ... which would make consistent sense in both annual and quarterly or monthly contexts

This space intentionally left blank

Technical spotlight: an end to confusion over incremental modelling of activity costs

Robin Bailey – Managing Director, ‘Mr STEM’

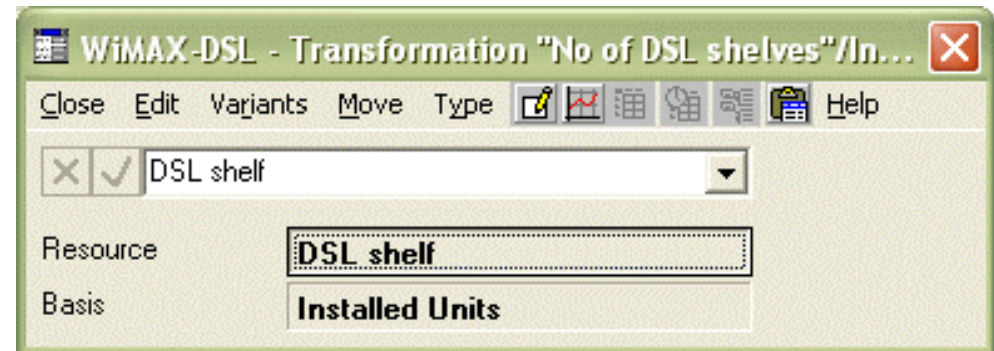
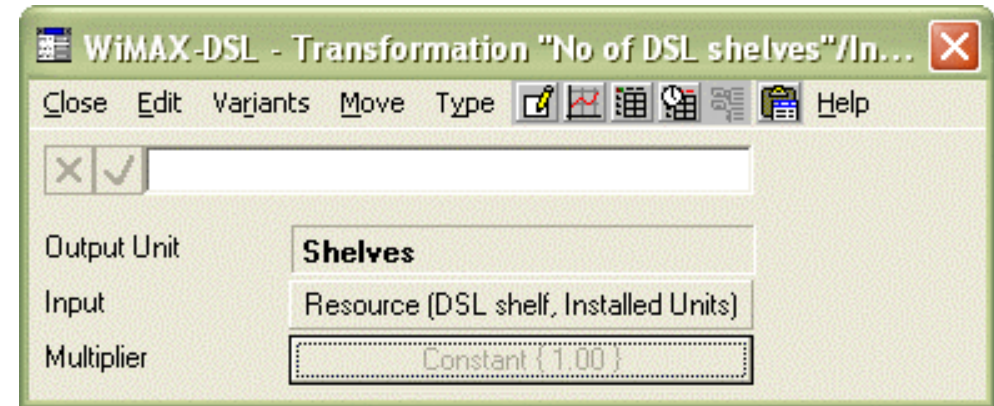


Introduction

- In the earlier *Opex work-out* session, we finished with two tentative ideas:
 - an *aggregate-mode transformation* (or output basis):
 - whose output would represent incremental demand ‘events’ within a period *only* (i.e., not cumulated over the year), and
 - which would be consolidated by summation (vs last period)
 - an aggregate-mode or *transient resource* whose costs would be associated with aggregate demand ‘events’ in the current period only, and not scaled as a proportion of a year
- So now we will examine these proposals in more detail:
 - examining a potential implementation design
 - working through some typical examples to test viability
- We will also mention some other current and directly-related development proposals in passing

Essence of a transformation


- A transformation was originally characterised in terms of:
 - its output unit*
 - e.g., Shelves
 - an input source element
 - e.g., Resource *DSL shelf*
 - an input basis
 - e.g., Resource Installed Units
 - a function [of inputs]
 - e.g., Multiplier
- We propose to add an Output Basis:**
 - *Aggregate or Instantaneous*
 - typically inferred from input basis
 - configurable in exceptional cases



- We need to think through the different transformation types to see what kind of user-interface (if any) is required for the output basis

Precise terminology is critical to minimise confusion

- ‘Basis’ could mean a *measure*, or an *attribute* – or a *mode*, or a *method*
 - An input or requirement basis defines:
 - a *measure* or *metric* of demand, or
 - a results *attribute* of an element
 - It is a choice of various different results available for an element, e.g.:
 - Resource Installed Units
 - Resource Incremental Units
 - Resource Used Capacity
 - An output basis defines:
 - a mode or method of interpretation; specifically
 - how to do consolidation
 - Assuming no attempt to fabricate complementary aggregate and instantaneous outputs, then there will continue to be *only one output number*
 - No choice of different output measures



- All input *bases* map to two output (or consolidation) *modes* – see over ...

Classification of input bases

Aggregate

- Resource (non-transient):
 - Incremental Units
 - Incremental Capacity
 - Capital Expenditure
- Service:
 - Traffic [volume in period]
 - Revenue [in period]
- Transformation:
 - Aggregate Output
- Resource (transient):*
 - [Incremental] Units
 - [Incremental] Capacity

Instantaneous

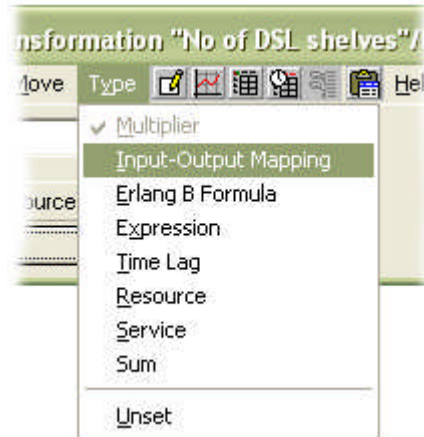
- Resource (non-transient):
 - Installed Units
 - Installed Capacity
 - Used Capacity
 - Utilisation
- Service:
 - Connections
 - Average Connections
 - Busy-Hour Traffic
 - Annual Traffic [rate fudge*]
 - Annual Revenue [rate fudge*]
- Transformation:
 - Instantaneous Output



Single-input transformations

■ There are currently six types:

- Multiplier
- Input-Output Mapping
- Erlang B Formula
- Time Lag
- Resource
- Service



■ But we are considering including a Multiplier input for all types:

- to streamline common modelling constructs
- e.g., Erlangs -> Channels -> Mbit/s

■ Resource and Service ‘identity’ transformations could be withdrawn

■ For each these types, the output basis would be inferred from the input:*

- with no explicit user-interface

■ You could potentially choose a specific output basis for requirements, but:

- only one Output result (Aggregate or Instantaneous) would be defined
- with sum and last consolidation respectively

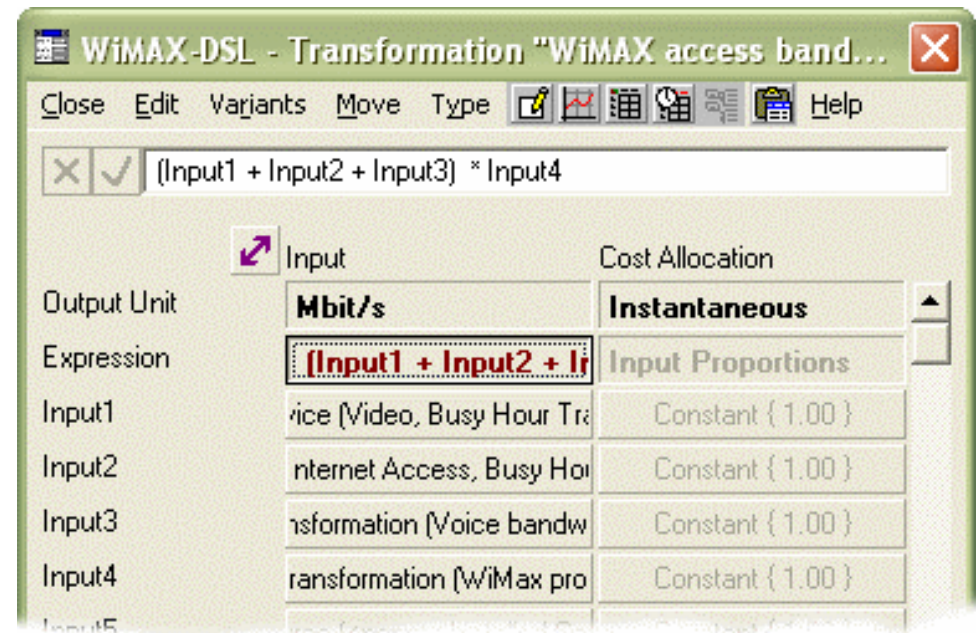
■ For ease of checking and backwards compatibility, the Output result would be linked to whichever of these were defined, but with consolidation = None

■ There would be no general, consistent way to fabricate the ‘other’ result as a generic delta or sum of the main result

* As per the previous slide

Multiple-input transformations

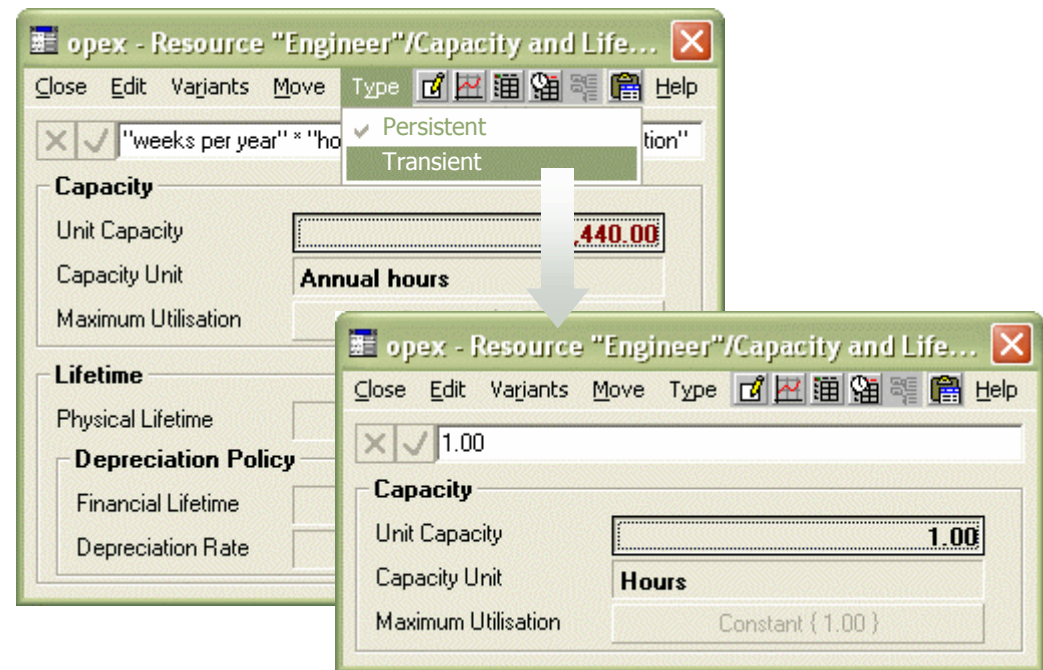
- A Sum should generate an error: at least if the inferred output bases do not match across all defined inputs, or even if the input bases did not match:
 - so treat like the single-input case
- In contrast, an Expression could quite reasonably combine input bases, e.g.:
 - different cost-allocation bases
 - scaling a sum of aggregate values by an instantaneous proportion
- So the output basis would require an explicit user-interface:
 - unambiguous and more flexible
 - even if any one aggregate input might require an aggregate output
 - a warning could be generated if this tentative rule were violated*
- Only one Output result (Aggregate or Instantaneous) would be defined, just like the single-input case
- Chart compatibility (consolidation and format) is likely if unit labels match, e.g.:
 - GBytes \Rightarrow Aggregate (sum)
 - Mbit/s \Rightarrow Instantaneous (last)



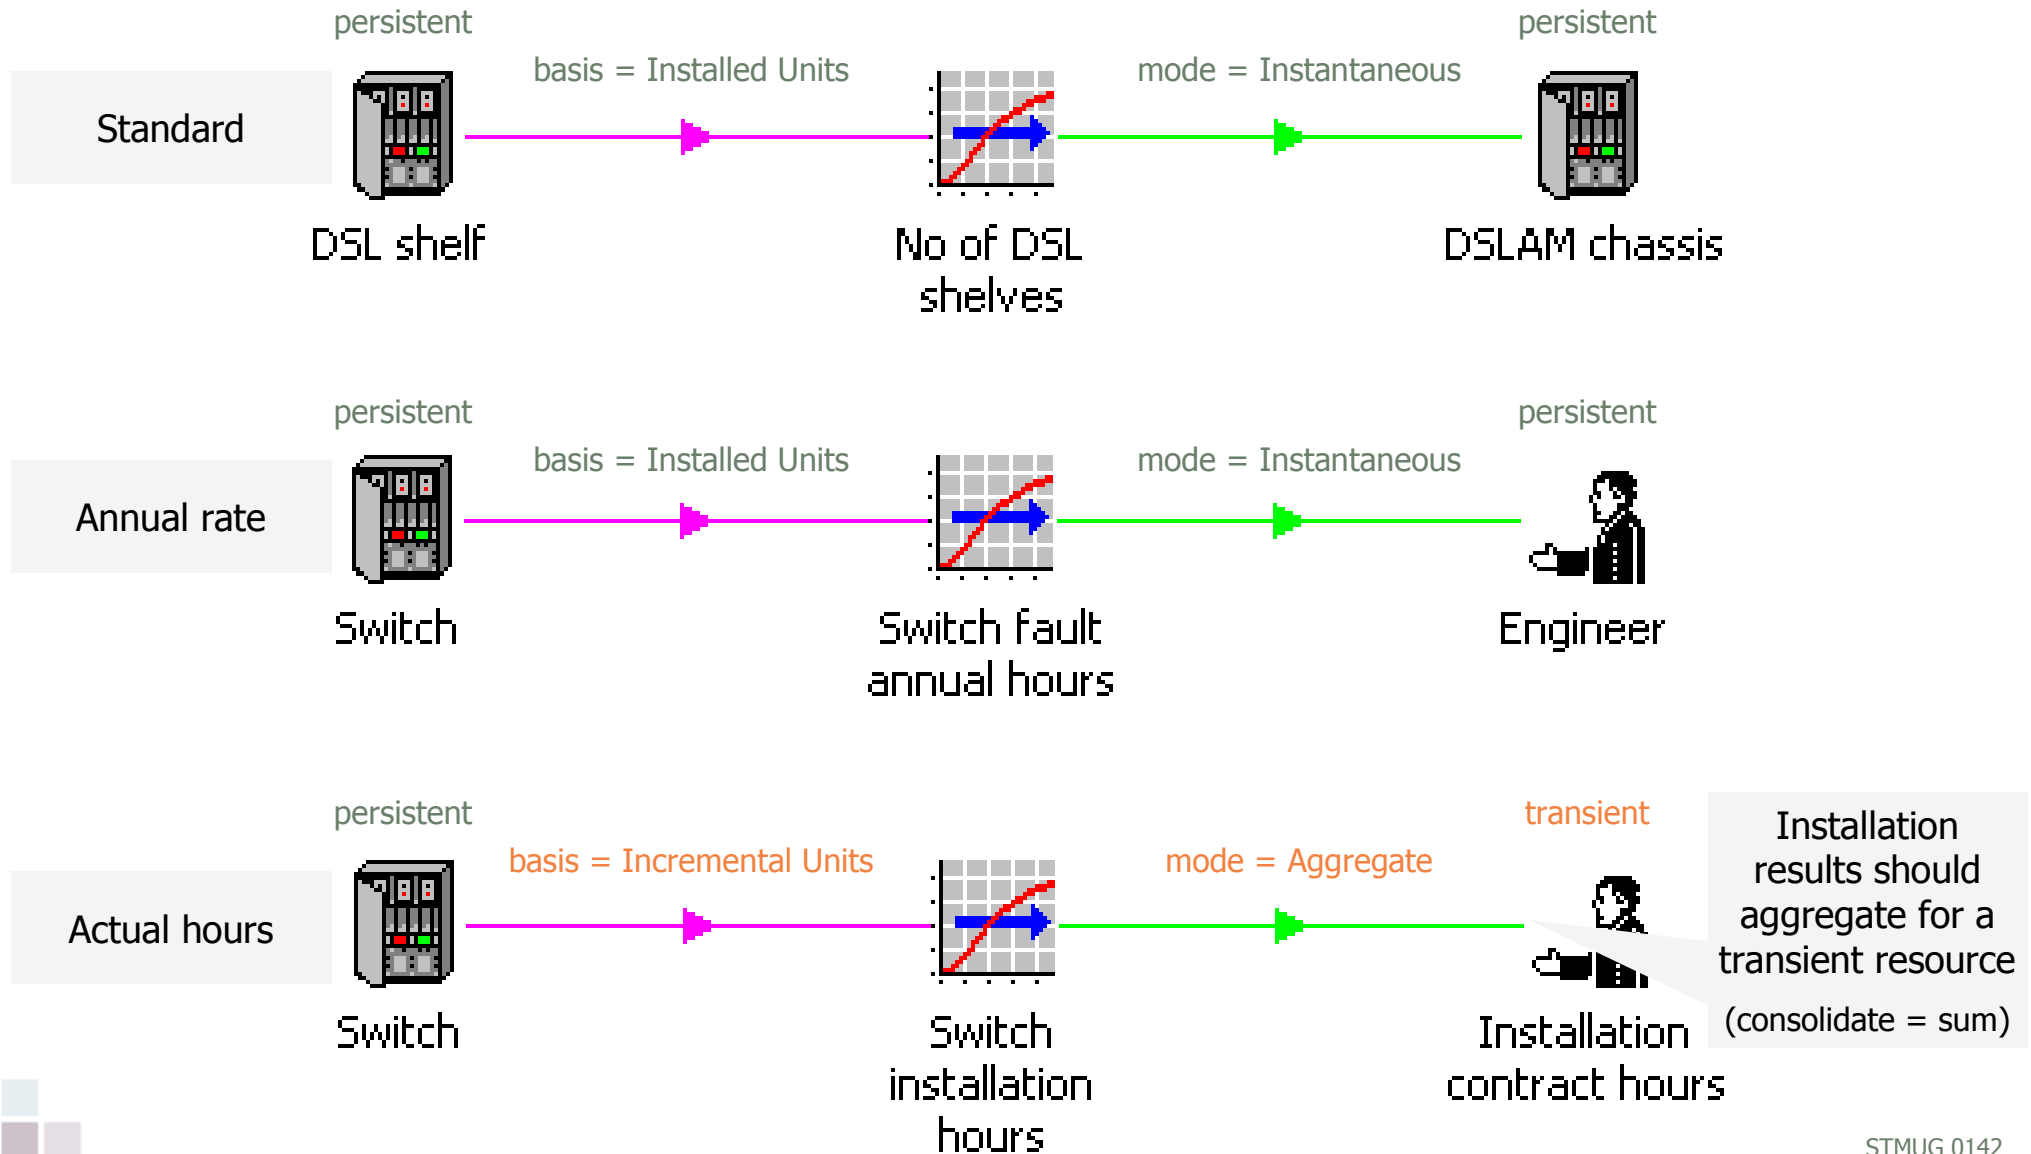
Persistent and transient Resources

- A resource could be characterised as *transient* if its lifetime were zero ...
- ... or via an explicit Life Mode* field ...
 - *Persistent / Transient*
 - and an error if lifetime explicitly set contrarily; i.e., zero / non-zero
 - or capital cost explicitly set for a transient resource
- ... or a *Type* menu could switch from persistent to transient mode:
 - which would hide the Lifetime inputs group, including the redundant Depreciation Policy
 - though the type menu is often hard for trainees to grasp
 - less clear than a visible drop-down

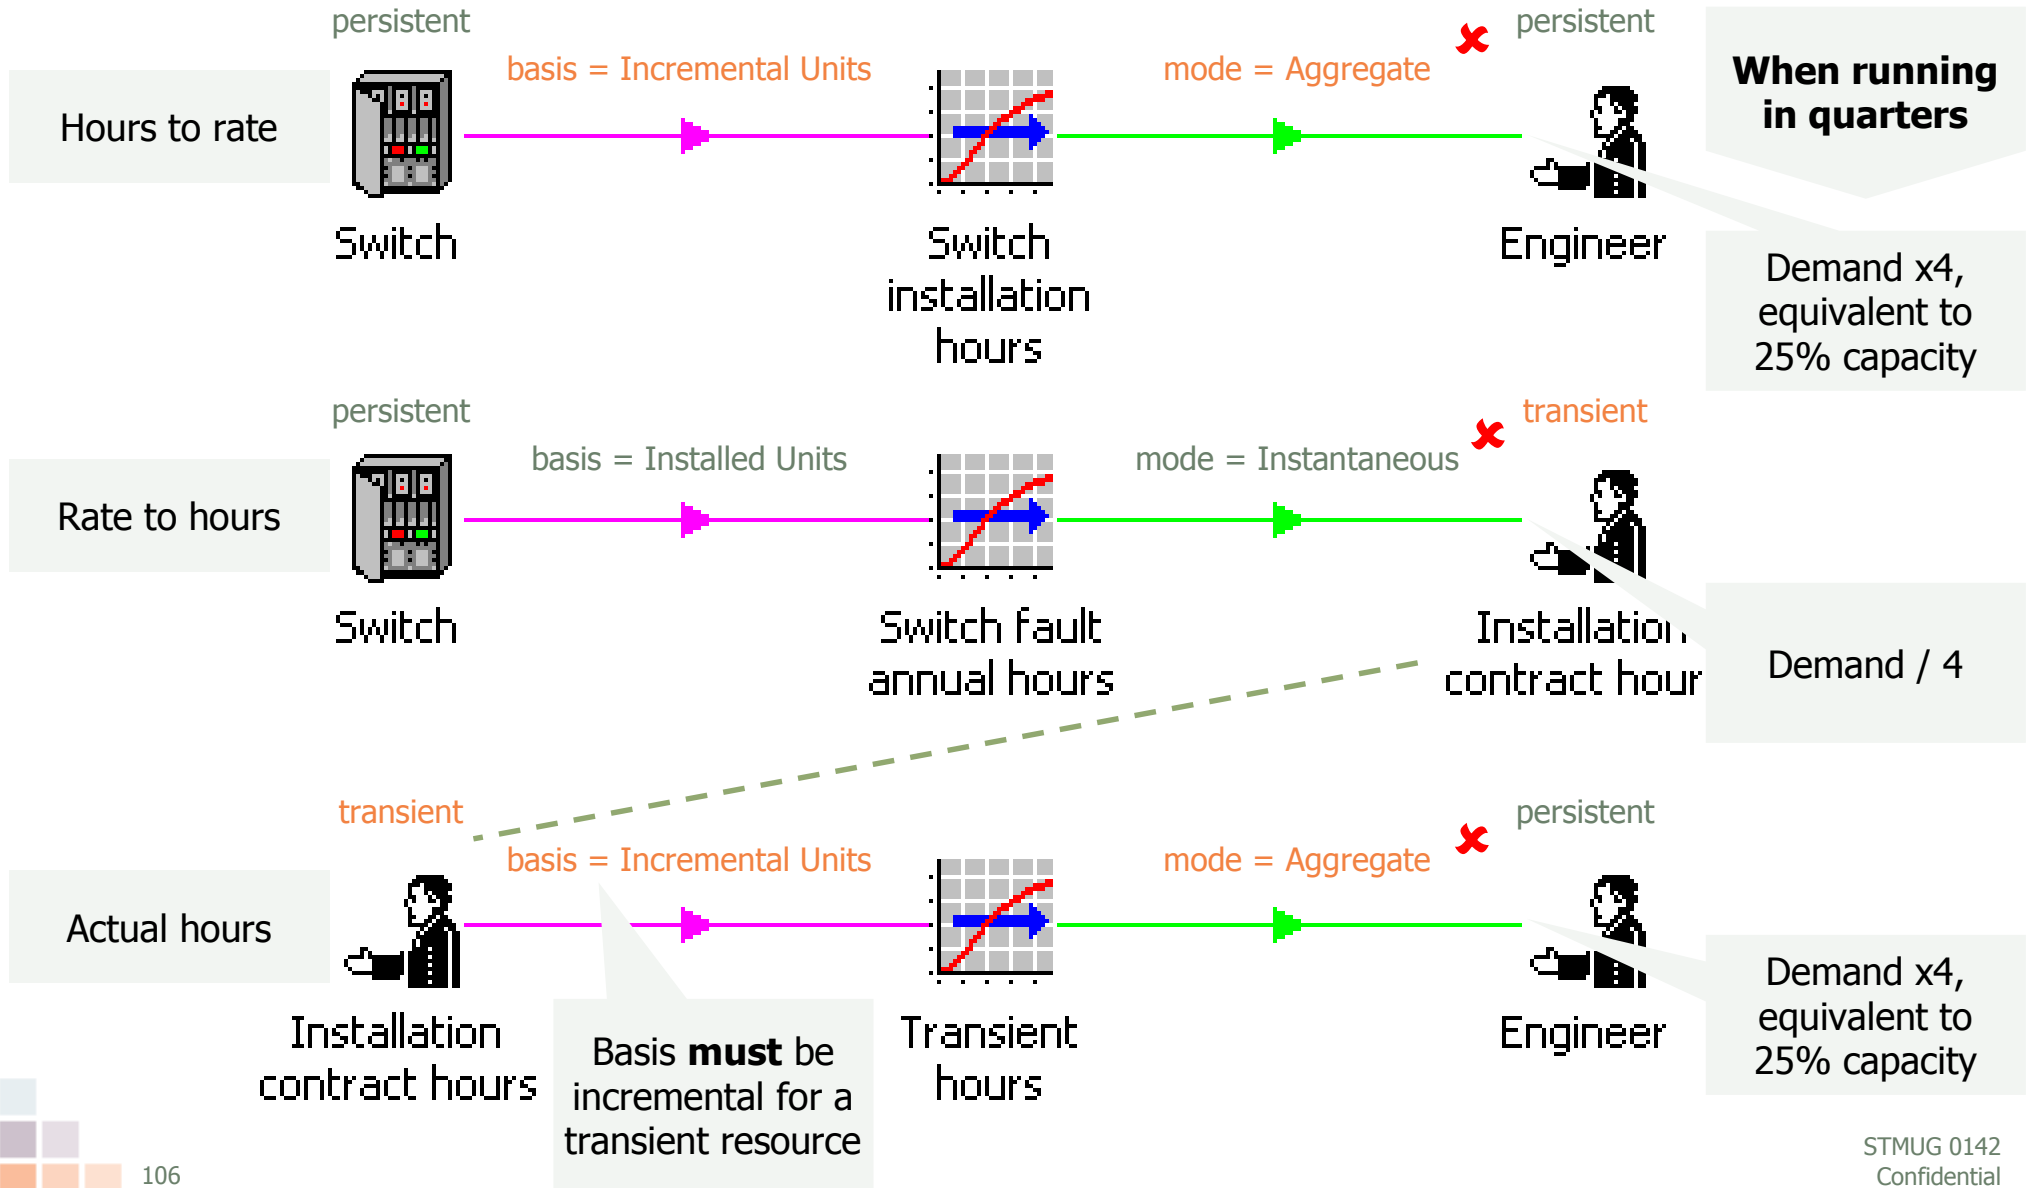
- Operating costs would be generated:
 - in proportion to the period length for a persistent resource; e.g., per engineer per annum
 - simply per unit capacity for a transient resource; e.g. per hour



Some contrasting examples



Switching between aggregate hours and permanent staff



Shorter lifetimes and shorter time periods

- If you do choose to map task hours onto permanent staff capacity ...
- ... then you may worry about retaining redundant staff beyond an initial installation surge
- A one-year lifetime is a long time if you are modelling in quarters or months ...
- ... so we may extend Lifetime detail:
 - additional Months input
 - lifetime = Years + Months
 - transient (or error) if both zero
 - warn if run with insufficient granularity; e.g., years when lifetime define in months
 - represents a fixed-term contract rather than a notice period

The screenshot shows a software window titled "opex - Resource 'Engineer' / Capacity and Life...". The window contains several input fields and labels:

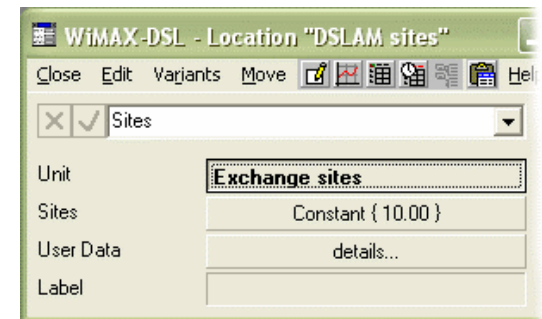
- Capacity**
 - Unit Capacity: 1,440.00
 - Capacity Unit: Annual hours
 - Maximum Utilisation: Constant { 1.00 }
- Lifetime**
 - Lifetime Years: 0
 - Lifetime Months: 3
- Depreciation Policy**
 - Financial Lifetime: 0
 - Depreciation Rate: -1.00

- This may be a lot of pain for little gain
- Cost per hour may be 'good enough'

Consistent and prominent placement of unit labels

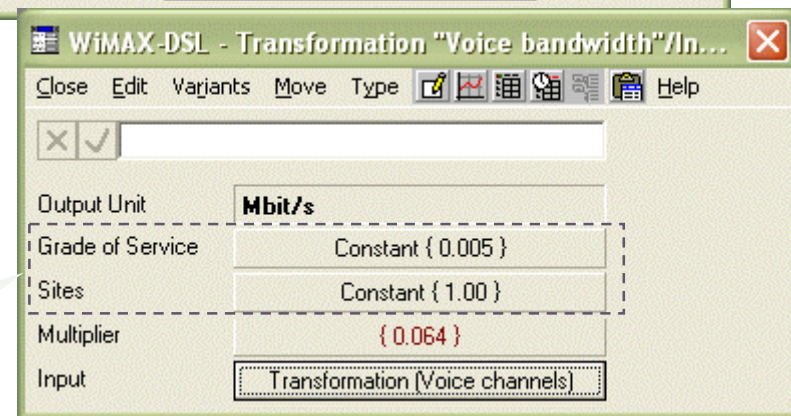
- Labels like *Connections*, *Call Minutes*, *GBytes*, *Erlangs* and *Mbit/s* have pride of place in the Service Demand dialog
- In contrast, Capacity Unit comes *after* Unit Capacity in the Resource Capacity and Lifetime dialog:
 - suggest switching the order and renaming Unit Capacity to avoid the mirror term (and everything is implicitly per unit anyway):
 - Capacity Unit
 - Physical Capacity
- Output Unit is buried in the Transformation Other Details dialog:
 - suggest moving the Output Unit to the top of each of the alternative Input and Transformation dialogs

- The unit labels of Market Segments and Locations are similarly buried ...
- ... and we have been planning to have ‘Size’/‘Sites’ in the icon menu lead to ‘the time-series button’ anyway:
 - suggest create new child dialog:
 - Unit
 - Size / Sites
 - or go straight to a re-ordered ‘top-level’ dialog:
 - Unit
 - Size / Sites
 - User Data
 - Label*
- and remove ‘Other Details’ from the icon menu



A multiplier for all transformations

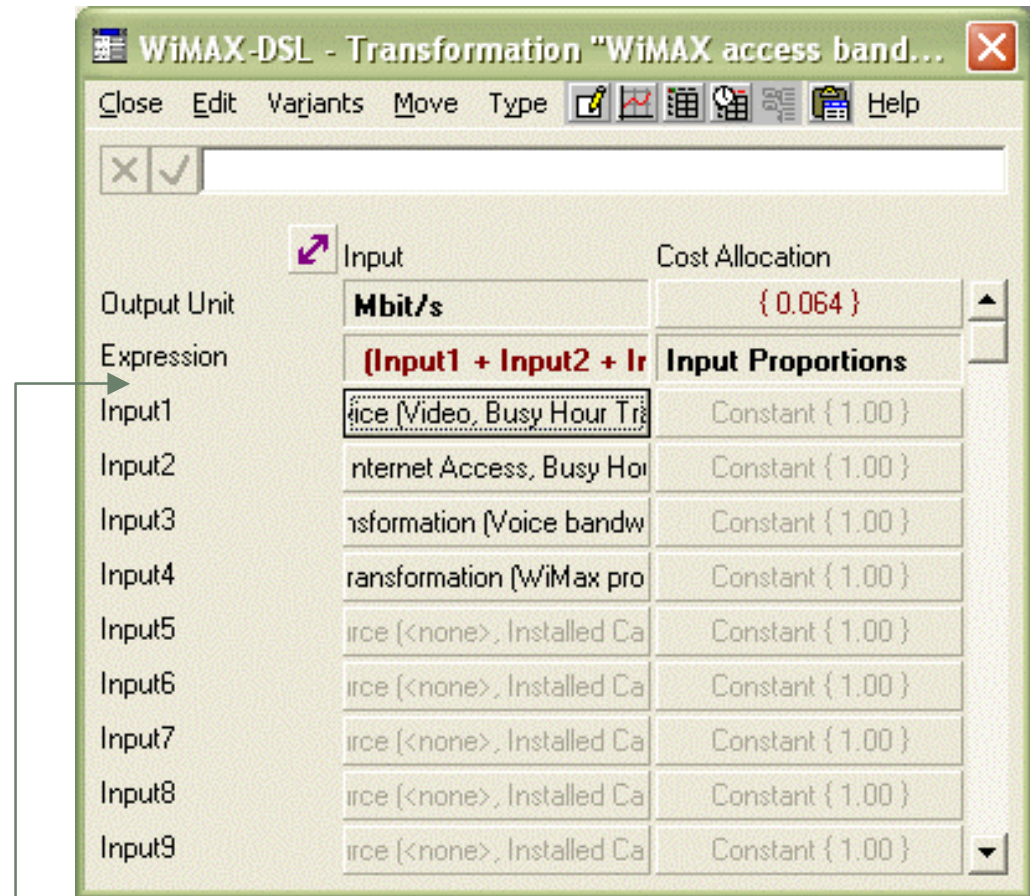
- As mentioned near the beginning, we are considering including a Multiplier input for all transformation types:
 - to streamline common modelling constructs
 - e.g., Erlangs -> Channels -> Mbit/s
 - also means the Multiplier input could persist if an expression transformation is required by connecting up a second input
- Like the Output Unit, the Multiplier input should appear in each of the alt. Input and Transformation dialogs:
 - keeps all factors visible
 - preserves backwards compatibility for Multipliers linked from Excel



It might be easier to read with a group label, **Erlang B**, around Grade of Service and Sites

But the expression transformation dialog gets messier

- It still makes sense to have a separate Multiplier input for an Expression transformation:
 - otherwise you have to bury one within User Data
- So we are looking at adding:
 - Output Unit
 - Multiplier
 - and possibly Output Basis*
- The existing dialog is already crowded, with the Cost Allocation basis in the same row as the Expression input
- Best suggestion is shown to the right:
 - 2-column grid doesn't really work
 - no idea where to put Output Basis




- Multiplier should follow Expression, for consistency, but would then be very easily confused with the Inputs**



Modelling clinic: live Q&A on current user roadblocks and some recent case histories

Robin Bailey – Managing Director, ‘Mr STEM’



Introduction

- This session is intended principally as a modelling clinic ...
- ... meaning an opportunity to take questions and feedback from the audience on an unscripted basis
- However, past experience suggests that:
 - either no one has any problems with STEM and everything is clear
 - those people have already left for the airport
 - or people are simply not prepared or comfortable to discuss live projects in this open forum
- Whichever of these applies, previous sessions on this theme have been embraced by the audience with a deafening silence ...
- ... so we have prepared a few short support case-histories which we can relate to provide a little insight into this vital aspect of the STEM experience and which may stimulate some input from the audience too

Context from a few STEM interactions worldwide

- An operator with headquarters in the Middle East is using STEM to understand the allocation of the costs of its radio and core network between different customer groups and service types (A)
- A leading wholesale operator in the Asia-Pacific region licensed STEM to create and test detailed financial models of a planned national broadband infrastructure, together with a range of scenarios for its future evolution linked to suitable customer and traffic forecasts (B)
- A US manufacturer of high-performance networking equipment works with STEM to create generic, client-focused models of its communications solutions which are then customised from time-to-time in direct engagements with individual customers around the world (C)
- Note: these letters in parentheses are not in any way linked to the identity of the respective clients. For the avoidance of doubt, none of these are the first letters of the companies in question!

A Extremely broad lookup problem (A)

- An operator with headquarters in the Middle East is using STEM to understand the allocation of the costs of its radio and core network between different customer groups and service types:
 - they have modelled all the physical infrastructure, as well as many layers of detail about the internal company staff process structure
 - the radio-network calculations are replicated 16 times across a combination of four geo-types, *Dense Urban*, *Urban*, *Rural* and *Highways*, and four cities or regions (which shall remain nameless)
 - cost-breakdown is enabled, meaning that STEM is reporting allocated costs for each of the ~ 200 services (and some replicated) down to the level of from which of the ~160 resources (and some replicated)
 - we calculated that there are in excess of 50,000 cost allocation nodes
- The model runs reasonably quickly, but the element-matching algorithm in the Results program had $o(n^{5-7})$ complexity, depending on how you count
- By careful analysis of the code paths, we were able to optimise the search strategy down to $o(n^{3-5})$, with a major impact on performance



B Lessons learned from a prolonged on-site client interaction (B)

- A leading wholesale operator in the Asia-Pacific region licensed STEM to create and test detailed financial models of a planned national broadband infrastructure, together with a range of scenarios for its future evolution linked to suitable customer and traffic forecasts
- In a first project, we provided training and worked with the client to develop an initial modelling framework, including modules for FTTH, wireless and satellite broadband, which could be evaluated across a range of representative geo-types
- We returned after five months to deliver additional training for recent recruits to the finance team, and to provide direct, expert modelling advice as plans were prepared for a strategic investment deadline
- Because of the length of the two on-site interactions (seven weeks in total), this single client inspired a number of incremental developments of benefit to all users, in addition to one sneaky manoeuvre which massively accelerated the scenario-generation process at a critical stage:
 - Also Office 2010 missing reference. Also Win7 install. Also auto-macro.

B Side-stepping an avoidable problem of immense scale (B)

- The FTTH network will connect almost 90% of the target market:
 - so the associated design and build process has been de-constructed in terrifying detail
 - and the calculations are also replicated across 16 geotypes (different ones!) by density
- One particular cut of the civil-works:
 - into three almost identical sub-models, coupled with widespread use of companion resources (via transformation) to calculate capitalised installation costs,
 - led to an unexpected and significant slow-down ...
- ... just when a process of intensive scenario running for investors was getting underway!
- An in-depth review of the model yielded a more economical solution to the capitalised installation problem* ...
- ... as well as a recommendation to re-do the civil-works cut more efficiently using template replication
- However, at its present exaggerated scale, it was actually the geotype replication which was taking too long
- Because the investor scenarios were all driven out of Excel, we were able to re-implement the template variants as STEM scenario variants ...
- ... and then use a hastily engineered results extension to perform a built-in 'sum over all scenarios' to calculate total revenue and investment numbers



C A little bit of self-help goes a long way

- A US manufacturer of high-performance networking equipment works with STEM to create client-focused models of its communications solutions
- This involves the maintenance of a number of generic solution models which are then customised from time-to-time in direct interactions with individual operator
- We have supported this client in their endeavour for a number of years:
 - our availability is well understood
 - we can usually read pretty well how urgent a support request is
- So ... we just missed a support email which arrived one Friday evening:
 - East Cost is GMT-5
 - The message was evidently not urgent (otherwise my phone would have rung!)
 - The client simply couldn't figure out the Installed Units result for one particular resource in the model
 - We did not respond until Monday morning, at which point it emerged that that there was a missing file
 - So we were not able to run the model until Monday afternoon (3 days later), whereupon the answer turned out to be a trivial detail (forgotten maximum utilisation input)
 - Whereas the built-in Audit function would have found this directly
 - Our client acknowledged that this was a perfect, missed self-help opportunity which would have avoided some degree of personal humiliation too!

Focus on removing roadblocks!

- If there is time, and there wasn't already time at the end of the *Platform evolution* session yesterday, then we can spend some more time looking into proposed developments
- These ideas have all arisen directly from project experience in the last year:
 - some are in response to witnessing certain client frustrations
 - some are just fresh insights from working full-time as a user for a number of extended periods
- We are effecting a gradual transition towards a closer engagement with customers, collaborating on projects beyond initial training to help clients steer a sure path towards well-engineered modelling solutions
- Implied Logic is careful not to present itself as a consulting company:
 - to avoid any conflict of interest with potential consulting partners
- However, our expert advisory services are adding significant value for operator, vendor and consulting clients alike

Robin Bailey

Managing Director, 'Mr STEM'

Mobile: +44 7776 198458

robin.bailey@impliedlogic.com

*Business- modelling software, training and expert advisory services
for strategy- planning and cost- allocation managers*

- Telecommunications
- Information technology
- Cloud computing
- Serviced properties
- Transport logistics
- Energy supply



This space intentionally left blank