ARTIFICIAL INTELLIGENCE AND DSS

Also in this Issue:

- Augmented Reality in Business
- Email (and PDFs) from Hold Files
- Pushing the Limits with OpenQM

intl-spectrum.com

INTERNATIONAL

Spectrum ®

36TH ANNUAL
MULTIVALUE CONFERENCE

LITCHFIELD PARK
ARIZONA

The Premier MultiValue User Event

MAY 1 – 4, 2017

THE WIGWAM
LITCHFIELD PARK, ARIZONA

www.intl-spectrum.com/conference

Register by March 22nd and Save!

The **International Spectrum MultiValue Conference** is the only place that provides diverse solutions and support for integrating and developing your Enterprise applications into the world of business and technology.

This annual event is the place to find the **knowledge**, **people**, **resources**, and **technologies** for enhancing your MultiValue (D3, UniVerse, UniData, Reality, jBase, Caché, QM, and OpenInsight) software applications.

INTERNATIONAL Spectrum

THE MULTIVALUE TECHNOLOGY MAGAZINE

DEPARTMENTS

From the Inside

Let's say you are a rock star or a professional athlete. You have talent. More than that, you are a master of your craft. Performing — onstage or on the playing field — is the practical demonstration of what you can do. Despite the hours you spend playing in public, you are still required to practice, practice, practice in private.

Let's say you are a programmer. Do you think you are a good coder? Ask yourself: Have I been practicing coding, or have I just been doing my job? Most of us assume that our day-to-day software development is our practice. But this is not the case. If day-to-day software development is like an NBA game – because that's where we make our money – it is clearly showtime and not practice. Basketball players use the time between games to improve their jump shots, layups, and passing. They do this for the purpose of increasing their skills and performance during the games. This is what you, as developers, also need to do.

One way to do this is to perform Code Katas. The term was coined by Dave Thomas, co-author of the book The "*Pragmatic Programmer.*" A Code Kata is an exercise in programming which helps you hone your skills. Although they were not designed with MultiValue in mind, they are a good fit. Let's go through one together.

Here's Dave's first Code Kata (for pricing goods at supermarkets) found at: http://codekata.pragprog.com/2007/01/code_kata_one_s.html

Kata One – Supermarket Pricing

Some things in supermarkets have simple prices: this can of beans costs $0.65. Other things have more complex prices. For example:

- Three for a dollar (so what's the price if I buy four or five?)
- $1.99 per pound (so what do four ounces cost?)
- Buy two, get one free (so does the third item have a price?)

This Kata involves no coding. The exercise is to experiment with various models for representing money and prices that are flexible enough to deal with these (and other) pricing schemes, and at the same time are generally usable (at the checkout, for stock management, order entry, and so on). It would be beneficial to us if we spent some time considering issues such as:

- Does fractional money exist?
- When (if ever) does rounding take place?
- How do you keep an audit trail of pricing decisions (and do you need to)?
- Are costs and prices the same class of thing?
- If a shelf of 100 cans is priced using "buy two, get one free," how do you value the stock?

More often than not, the requirements of our day job won't parallel our needs when it comes to growing our skills for the future. The way to get ahead of the curve, and to keep our minds fresh, is to look for side projects, for practice exercises, or any other excuse to stretch and improve. Do that, and *you* are a rock star.

NATHAN RECTOR
President
International Spectrum
nathan@intl-spectrum.com

MISSION STATEMENT *International Spectrum* magazine's editorial mission is to be the premier independent source of useful information for users, developers, and resellers of MultiValue database management systems, open systems business database solutions, and related hardware, software, and peripherals. Published bimonthly, *International Spectrum* provides comprehensive coverage of the products, companies, and trends that shape the MultiValue marketplace as well as the computer industry at large — helping its readers get the most out of their business computer systems.

twitter.com/intlspectrum

intl-spectrum.com/facebook

GET CONNECTED.

KNOWLEDGE AND EDUCATION FOR THE MULTIVALUE PROFESSIONAL.

ABOUT OUR PROFESSIONAL MEMBERSHIP

We are all busy in our day-to-day work and staying up-to-date with the current MultiValue technologies can be difficult.

Professional Memberships provide you access to knowledge, solutions, information, and code that you won't find in other locations.

Professional Membership Includes:

- Magazine in Electronic and Print Formats
- Newsletter
- On-Demand Videos
- Live Webinars
- Discounted Conference Rates
- Research papers
- Case Studies
- Source Code

INTERNATIONAL
Spectrum ®

http://www.intl-spectrum.com/membership/

AI and DSS

Part 1

BY CHARLES BAROUCH

In science fiction, we see all sorts of Artificial Intelligences (AI) who seem to want to be humans or want to replace humans. I've met humans. If that's their goal, the intelligence part might not be that strong. We're going to take a medium-deep dive into AI and Decision Support Systems (DSS). Most of it won't be language or database specific, but all of it can be done in MultiValue.

Smarter Animals

One of the basic AI examples is a very simplistic learning program called Animals. I like this example because it turns to ruin relatively quickly and that's something you need to understand about AI, especially when it is used for DSS. The odds of doing it wrong are surprisingly high. This is just one reason I can't see a computer AI wanting to be human.

This programming idea has long roots. For example, here's a link to the Univac assembler version of the program: http://www.fourmilab.ch/documents/univac/animalsrc.html.

Spoiler alert, nothing is simple.

The code for Animals is easy to create in any language. Here it is in Perl http://www.perlmonks.org/?node_id=195991. Here's the vanilla MultI-Value code so that you can try it yourself https://www.intl-spectrum.com/resource/781/default.aspx.

Spend a little time playing Animals and you'll notice that pig, for example, might end up on multiple paths. While that does mimic the real world *(i.e. there's more than one way to guess your way to pig)*, it makes for a messy database very quickly. Can we code some synergy back in to fix that? Yes. Yes we can. Can that lead to asking duplicate questions? Yes. As you play it, you'll see the flaws. This is a hands-on exercise. Can we code around the duplicates? Yes. Indexes, combined with some fuzzy compares ("Does it has claws?" and "Do they have claws?" are the same question, for example; we

need an AI smart enough to see that) can solve this. And that will create the next problem: too many dead ends. And so it goes.

If creating a simple game like this looks like an endless spiral descent, how can we ever hope to create something more complicated?

Combined Pricing

In one of my previous lives, I consulted for a company that made additives for plastic. What does that mean? It means that they owned the formula for making Tide bottles the right shade of orange. It means that when you smell the Johnsons & Johnsons' Baby Powder, you are smelling their additive — the bottle is sealed — you aren't actually smelling the contents. This was a company who sold their product with a unit of measure based on the size of commercial railway shipping cars.

The problem they had was pricing. Shipping a half-car's worth of midnight black additive was just as expensive as shipping as a full-car, so they

would discount other products to get you up to full-car orders. This meant that if you ordered a half-car of rust red additive and then filled the remainder with a quarter car's worth of cobalt blue additive and a quarter of midnight black, you might get a different price than if you started with cobalt blue then added the other two with the 'fill the car' discount. Worse, different sales staff worked the calculations differently. There was wiggle room, so that the computer would accept the negotiated prices.

The company wanted a DSS where you could fill-in the entire order and have it optimize the discounts, either favoring them or favoring the client, based on the relationship. Slow payer, small buys, and new customers got lesser deals while the biggest and most loyal got the best discounts.

For this project I needed an algorithm which gave an array of valid answers and then picked one based on the profile. If you think about it, there's really no AI here. This is not intelligence because it should not learn in the AI sense. All it needs to do is brute force: calculate every possible deal. In the midnight/rust/cobalt example, there are only six solutions: (1) midnight then rust then cobalt, (2) midnight then cobalt then rust, (3) rust then cobalt then midnight, (4) rust then midnight then cobalt, (5) cobalt then midnight then rust, (6) cobalt then rust then midnight. Once we have the six solutions, we rank the buyer and give them the one they qualify for. This is DSS but it really isn't AI. At the end of the day, combined pricing was easier than coding animals.

Buying a Hard Drive

Let's take a better example. This is a simple one. *Spoiler alert, nothing is simple.* I want to buy a hard drive from the local computer store. I need it for my laptop and I want one terabyte. The query equivalent of my question might look something like: SELECT INVENTORY IF TYPE = "HARD-DRIVE" AND IF FORMFACTOR = "2.5" AND IF CAPACITY = "1TB." How is the query different than the question? The query will select all the possible drives but I'm only looking to buy one. The missing element is analysis. That's a big difference. No one, statistically speaking, wants to buy one of each. This is where AI can serve us well.

For example, we could take the hundred and six drives that the query returned and put them up as a web page. When the requestor does a compare, we do an analysis of why they might have been drawn to those drives. If they were the cheapest, we know to re-organize the results by ascending price next time this requester looks for drives. If they were the fastest, we could work with that, too. The more he or she pokes around, the more data you gather. Likewise, we weigh-in the purchases, not just the compares.

A million customers later, you can average peoples' experiences to make some assumptions about new buyers based on the previous ones. Then we can layer on the Amazon-style people-who-bought-this logic. Of course, the query itself needs a sprinkling of AI. I said one terabyte but a good AI would know that nine-fifty meg drives might be worth listing. Of course, up-selling is good too, which would put a few higher capacity drives on the list, just like a human salesperson would. With a little effort, we could make a smart laptop drive recommendation system. That's not what we want. It needs to be more general. That, in turn, complexifies things. Desktop, NAS, and server drives need different rules than laptop drives. For example, RAID might be an issue in these less portable devices. FORMFACTOR is more than just a criteria; it is the key to deciding which secondary rules should be considered.

Likewise, it is easy to fall into false equivalences. Memory also has speeds, but hard drives measure speeds in RPMs where RAM doesn't. Computer enclosures also have capacity but measure size in bays and slots, not bytes.

The reason we don't just want to make an assumption about what to present and how to order it is simple: lost sales. When Joe is adding a drive to his personal machine, he might be looser on price but really care about speed. Sandra is all about MTBF (Mean time

between failures) because she's building mission-critical servers. Jenny, on the other hand, is cranking out three to five machines a day, so even a dollar difference in the price of a component is something that will add up rapidly when you look at the thousand or more machines she'll build this year.

A good sales associate knows that one-size does not fit all, neither in product nor in sales approach. If our AI is really good, it won't just reduce selections (tailoring) and optimize the order it presents them in, it will also see the bulk/frequent buyer, like Jenny, differently than it will Joe. Jenny should be offered loyalty discounts and price structures that encourage her to stock up instead of buying as-needed. You might even offer her discounts if she puts your logo on her marketing materials.

You can see that AI isn't just a gimmick, like faster sorting or smarter selects. The thing which makes it AI is the adaptive nature. AI learns. And to write AIs, we must also learn.

Next issue we'll talk about the dystopian fiction of evil AIs and show you how MultiValue envy is the secret weapon destined to keep humans on top. **IS**
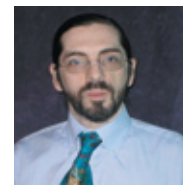
CHARLES BAROUCH is the CTO of HDWP, Inc. and the Publisher at HDWPbooks. You can read his writing in *International Spectrum*, *Theme-Thology*, *Novo Pulp*, *Pax Solaria*, *PerehelionSF*, and the *Interrogative* series, which begins with *Tiago and the Masterless*.

Augmented Reality in Business

BY NATHAN RECTOR

As some of you know, I've been tracking the changes in AR (Augmented Reality) pretty closely. I've even been experimenting with it some, looking for those business cases that would explain why, when, and how it should be used.

For obvious reasons, AR games are being used to demonstrate this technology. Games are visual, and being able to play in a digital fantasy world is really cool! Because you're seeing it first in games, and not in business applications, it can obscure the other value propositions offered by AR.

Before we get in too deep, we should address a bit of a misconception: AR *is not* VR (Virtual Reality). There is a distinct difference between the two technologies, even if they seem similar. AR offers users graphical enhancements to their real, physical environment, rather than creating an entirely new environment.

VR is designed to replace what you are seeing with a completely new world. This is what Samsung shows in the commercials for the preview of the very popular television series 24. This what Oculus Rift offers when you explore the galaxy using their headset.

> *If you thought mobile phones and e-commerce changed the business world, wait until augmented reality really takes hold. And the only thing that is holding it back is the hardware.*

If you have ever played Pokemon Go in AR mode, then you have an idea of what AR does. If not, consider this: how would you like to look at a barcode and have a window pop-up in front of you and tell you that you have a quantity of ten, and that you have five orders, so you need to order twenty-five more? AR is all about layering information, or animation, on top of real world triggers like sight or location.

Some of the AR business cases are pretty clear, but we've barely gotten started. I don't know where this will end up, and how far we will push the technology, but I can make a few guesses. Some people think AR will be a fad. Remember Google Glass?

To understand why I don't expect it to be a fad, let's look at a bit history. After the desktop seemed to emerge as a clear winner, the tools of portable digital access started to carve out a place; cell phones, Palm Pilots, and PDAs. Then someone (Apple) legitimized the smartphone. The iPhone wasn't the first smartphone, but it was the one that breathed life into the category. Laptops and tablets helped erode the dominance of the desktop.

If you remember, when iPhone 1 was released, people wondered why it had a camera, and said it wouldn't replace the existing digital cameras. Many saw it as just another device to play games on. The smartphone went from toy to business device in a very short time.

As processing power increases, the device in our pockets will continue to take territory from the device on our desks. AR can be done with your smartphones (again, Pokemon Go and Snapshat's photo alterations software for example), and the only thing that is limiting it is how weird you look wearing the headsets. And that's changing, too.

Several companies already make fairly simple glasses that can project flat images for their wearers. They are increasingly popular in warehousing and manufacturing firms, who can use

them to issue instructions to employees while leaving their hands free.

Meanwhile, companies such as Magic Leap, Meta, and Microsoft are building much more capable headsets that can sense their surroundings and react to them, projecting convincing, three-dimensional illusions onto the world. Microsoft is already running trials of its HoloLens headset in medical schools, giving students virtual cadavers to dissect, and architectural practices where several designers can work together on a digital representation of a building.

Designing a nifty piece of technology, though, is not the same as ushering in a revolution. Social factors often govern the path to mass adoption, and for AR, two problems stand out. One is aesthetic. The HoloLens is an impressive machine, but few would mistake it for a fashion item. Its appearance makes its wearers look creepy and disconnected from the real world. It needs to look "tech-geek cool" instead. One reason the iPhone was so successful was that it was the beautiful design.

The other big problem: Consent. Look at the major issues that Google had with Google Glass. While it was the closest AR design to tech-geek cool, it had a lot of social issues attached to it. You know that "camera and microphone always recording" thing.

Smart TVs that are always-on, and the over-sharing on Facebook, Twitter, and YouTube are part of the process of creating a new social normal. One that will no longer hinder AR headsets.

These problems aren't something that would keep a business from using AR headsets internally. As current makers of AR headsets make their hardware lighter and easier to use, it will be more common for businesses to use them in-house. So the question isn't why would I look at this tech for business, but how should I use it.

Inventory Control and Warehousing

Right now, we use PDAs, smartphones, or other devices to look at and update our inventory. But this is all done one at a time. Users have to check each item on the shelf individually. Augmented Reality brings all this information together. It can overlay the information onto the shelf so that you take it in all at once. If you go a step further, you can then overlay alarms and highlighted sections of the shelf where the computer algorithm identifies problems.

Once you know where on that single shelf there might be a problem, you can step back and look down the aisle, seeing all the items and shelves that have problems on them. Your employees no longer need to hunt but instead will have clear visual clues.

Training and Education

Manufacturing, service and repair, or just about any job has a large amount of on-the-job training. Most of that training is done by people who are simply retelling the story of their experiences. You can make training and education more effective by implementing in-context instructions or layered graphics. This allows employees to show complex processes step-by-step, or give prompts and instruction on a certain task, even while not even physically present.

Showcasing Products

The rise of e-commerce has made showcasing products especially important. With AR, online shoppers are able to fully inspect the product they're considering purchasing as if it

were really there. One popular use of AR for this purpose is software that allows potential buyers to place different three-dimensional models of furniture in their homes to see how the piece might fit in before buying.

Not only is the customer empowered during the decision-making process, but it also helps reduce future expenses to the business by minimizing returns.

Augmented Office Spaces

One potential offering of the augmented office is the ability to alter plain office spaces. Imagine a room where the only tangible things were people and load-bearing furniture, like tables and chairs. Beyond that, the walls are bare and windowless. Sounds kind of ugly, cold, and, well, unwelcome.

Augmented reality could actually transform that room for whatever purposes you require. Everyone has their own workspace requirements. Tim wants a huge window. Nora wants a dark room with plants. You like a two-monitor desk, but I need four monitors. You get the idea.

When digital information is overlaid on physical world objects, you can turn one of the office's plain walls into a window to the manufacturing floor of another city. And on the opposing wall you can have a shared whiteboard that can be used by a team of developers that work remotely.

Instead of opening windows on a monitor, you open windows into the physical world and generate a depth of windows. We aren't restricted by the number of monitors. We are only restricted by the physical and digital space needed to do the job.

Conclusion

Augmented reality won't exactly replace business software, but it will cause business software to present information in a totally different format. Just like charts and graphs create visualizations on columns and rows of numbers, AR will bring the boring business software into the new and fun.

If you thought mobile phones and e-commerce changed the business world, wait until augmented reality really takes hold, and the only thing that was holding it back was the hardware.

And the hardware is here now. **IS**

NATHAN RECTOR
President
International Spectrum
nathan@intl-spectrum.com

FROM THE PRESS ROOM

Paradigm Systems Releases Mercury Console 4.1.13

Paradigm Systems today announced the release of version 4.1.13 of its Mercury Database Console. This latest release resolved several open requests by clients to have improved sorting and quicker display of database files. Jay LaBonte, founder of Paradigm Systems and creator of Mercury Database Console, stated: "This release is the first in a series that will focus on specific client requests highlighting our commitment to provide the best U2 database management system available." ■

ONgroup Scales Up With MVON# Running in .NET

ONgroup Intl now helps organizations boost their value by running MultiValue applications directly in Microsoft .NET. Companies can run in the .NET or Mono CLR, the Common Language Runtime. MVON# is the toolset that allows sites to keep what works, while removing barriers to integration with contemporary technologies.

In addition to more obvious gains when using pervasive Microsoft technologies, this alignment with the industry alleviates the strains of interoperability. IT is able to move many tasks directly into the hands of end-users while other tasks go away completely. The ease of managing applications with satisfied users yields an overall increase in the value of those applications.

ONgroup Intl is scaling up with the formation of ONgroup Africa and the global distribution of MVON#. The idea of running MultiValue as managed code in the CLR without a MultiValue-only run machine originated in South Africa, where the MVON# toolset is already in use. "By decoupling applications from the DBMS, sites have the first horizontally scalable MV applications," said Grant Hart, Chief Architect for the MVON# toolset. MVON# is optimized to run with SQL Server and is architected to run with any MultiValue or SQL-based DBMS on the back-end, including Rocket U2 and Oracle.

Developers with their own application generator can use it as a .NET application generator. The opportunity to run in .NET is also extended to sites running in System Builder or SB+ who can experience the new browser user interface for their existing application using MVON# Netbuilder.

"We added the sharp symbol to the name to give a hint that with this toolset, MultiValue source code transpiles to C# to run in .NET. Running applications side by side with everything else running in the Common Language Runtime will be as obvious a choice in the coming years as running applications in common operating systems is today," said Charles Stauffer, CEO of ONgroup Intl. ■

Pick Cloud, Inc. Joins GOOGLE Cloud Compute Platform Technology Partner Program.

Pick Cloud, Inc. announced they have joined the Google Cloud Compute Platform Technology Partner program.

"We recognize that almost every product or IT service will either reside in the cloud or be cloud attached. For that reason, we have chosen to use Google exclusively. We are combining the power of Google's security and infrastructure with MultiValue databases and applications. The same data centers that power Google now power Pick Cloud and offer our customers the same security, stability, and availability. These can come at a high cost. Not with Google, now more than ever, they make it cost effective for new clients to consider a cloud solution.

Our customers are in the same company with the likes of Coca-Cola, Best Buy, Spotify, Motorola and many more, says Mark Pick, CEO and President of Pick Cloud, Inc. We also chose Google because they take security very seriously. Over five million organizations across the globe, including 64% of the Fortune 500, trust Google with their most valuable asset; their information."

This program includes partners such as systems integrators, IT consultant firms, developers, and other professional services companies that help their clients design, develop, and manage new cloud-based solutions. ■

Ladybridge Systems Releases 3.4-9 of QM

Release 3.4-9 of the QM Multivalue Database is now available for download.

This release corrects a bug introduced at release 3.4-8 where use of INS, INSERT() or <-1> to append a field to a dynamic array could cause a following field extraction to return incorrect data.

There are also a few minor enhancements as requested by our users.

The IF command in a paragraph now has an option to perform case insensitive relational comparisons.

A "B" option has been added to commands within the SHOW utility to move to the bottom of the available items.

An ECS qualifier has been added to the QMBasic OPENTEMP statement to create the file in ECS mode.

A PARAMS property has been added to the !CALLHTTP.CLS class module to provide an alternative way to pass URI parameters in a GET request. ∎

Rocket UniVerse Now Supports Python

Rocket Software announced that the Rocket® UniVerse application platform now includes support for the popular Python programming language. Other user-requested features include enhanced replication for high availability and disaster recovery, as well as streamlined audit logging to automate compliance documentation for regulations and standards such as HIPAA and PCI.

UniVerse is a fast, flexible database and application platform that mimics the way users think, accelerating the creation of data-driven applications in manufacturing, retail, and financial sectors, among others. Long known for its small administrative and IT footprint, UniVerse offers users lightning fast data access in an increasingly data-driven world.

The addition of native Python support opens UniVerse development to a new generation of programmers versed in open source language and tools, making it easier to hire qualified engineers and developers.

"Our customers have told us that attracting and retaining developers can be a challenge," stated John Bramley, VP of Application Platforms at Rocket Software. "We extended our application programming capabilities to streamline coding, take advantage of third-party tools, and tap into the large development community that is already fluent in Python."

Audit Logging helps organizations document compliance with tax and federal regulations by tracking all steps taken using UniVerse through a new streamlined audit architecture. In addition to aiding compliance with industry-specific regulations and standards, companies can use audit logs as a critical component of a risk management strategy to prevent potential problems before they occur and proactively identify potential fraud. Enhanced replication capabilities offer robust disaster recovery options to ensure organizations address increasing demands to protect critical information and prevent catastrophic data losses.

"With UniVerse and its enhanced replication, we have peace of mind for our support and development departments knowing that we don't have to spend unknown hours fixing unknown amounts of file corruption," said Ryan Odgers, IT Manager at OUTsurance Insurance. "We can now quickly and easily create a replica, saving us $30,000 a month in hardware support costs that we currently have to pay." ∎

Zumasys Releases jBASE 5.6

The newest release of jBASE 5.6 offers even more incredible features, such as native RESTful services, native data encryption, more granular audit logging and improved transaction journaling.

jBASE 5.6 includes over 30 enhancements and fixes, and also includes new features:

- **Native data encryption:** Get built-in data security and protection with the standard jBASE product for no additional cost.
- **Additional audit logging features:** Increase the granularity of logging options for deeper event audits. Logs can also be applied to a remote copy of jBASE to accelerate the backup process and reduce recovery point objectives.
- **Improved transaction journaling:** Quickly enable and maintain jBASE replication with new enhancements, including definable automatic purging. Transaction journals can be used to quickly bring older database copies up to date.
- **Updated System Manager:** Enable centralized management via a graphic web interface for encryption, audit logging, and transaction journaling.

Release Notes:

http://www.jbase.com/releasenotes/5.6.0/Release-Notes56.html ∎

FROM THE PRESS ROOM

Information Technology Systems Announces Comprehensive Support for Microsoft and Linux Infrastructure

Information Technology Systems is providing Microsoft and Linux Design, Support and Service to provide a single point of support for your entire network. The ITS infrastructure support team has experience supporting all aspects of networks including wide area networks.

Do you want to mix and match Windows and Linux services?

Do you want to add Email Continuity and Archiving to your Exchange?

Do you want to use Linux Apache web services?

Our support staff handles it all.

Services include support for:

• Windows (Server and Workstations)
• Exchange (add email continuity and archiving)
• Email continuity and archiving
• Sendmail (Linux email service)
• Apache (Linux Web service)

Go to http://www.itspick.com/infrastructure-support.html then contact us and let us put together a low-cost monthly support plan designed with you and your budget in mind.

Put the ITS Infrastructure Support team to work for you today, and experience support your way. ∎

Email from U2 Hold Files as PDFs

BY FELIX A STELLMAN

Many years ago, my client wanted to start emailing reports from our Uni-Verse systems. We had several ksh scripts that tried to build the headers for using Unix mail to send them out. Each time a new report was dreamed up, a new script was requested. After four or five of these that only sort-of worked, I started to look for a better way. One of the problems was that PCs had become a major tool and

...when emailed, the pdf report is much smaller than the original text report.

when employees would try to look at the reports on their PC it wasn't usually a pleasant experience.

Printing those emailed reports was even worse. I decided that we were missing a

valuable tool because we were ignoring the PCs and what they could do for the employee. While talking to one of my Unix admins, he came up with a tool to convert text into a PDF.

In Solaris there is postprint and in Linux there is enscript. I wrote a small script that my developers could call so that they could convert any text into a PDF *[Figure 1]*. This puts the printout into a form that can be easily viewed and printed by a PC without any spe-

```
#!/bin/ksh
# convert text file in hold file into pdf and send to portal
# $1 = file name
# $2 = source directory
# $3 = destination directory
# $4 = (L or P) source Landscape or Portrait
# $5 = remove
sysname=`uname -n`
remove="$5"
orientation="$4"
PATH=$PATH:/opt/sfw/bin
cd $2
if [ "$orientation" = "L" ]
then
# cat $1 | /usr/lib/lp/postscript/postprint -pl -m.8 -x-.03 -y.24 - > /tmp/$1
# cat $1 | /usr/bin/enscript -B --font=Courier8 -r -1 --margins=9:9:36:36 -o /tmp/$1
# changing to enscript codes to emulate what printers do fas1 10/20/10
cat $1 | /usr/bin/enscript -B --font=Courier8 -r -i1c -1 --margins=28:0:0:-24 -o /tmp/$1
else
# cat $1 | /usr/lib/lp/postscript/postprint -pp -m1.2 -x-.01 -y.24 - > /tmp/$1
cat $1 | /usr/bin/enscript -B --font=Courier11 -R --margins=18:18:9:9 -o /tmp/$1
fi
/usr/bin/ps2pdf /tmp/$1 /tmp/$1.pdf
cp -p /tmp/$1.pdf $3
rm /tmp/$1.pdf /tmp/$1
if [ "$remove" = "remove" ]
then
    rm -f $2/$1
fi
exit 0
```

Fig. 1 printout.converter

```
# /usr/lib/lp/postscript/postprint -pl -m.8 -x-.03 -y.24 - | lp -o nobanner -d top2
# awk -f /usr/local/bin/lttr_lnd66.awk | lp -d top2
/usr/bin/enscript -hBZ --font=Courier8 -r -i1c -1 --margins=10:0:0:-15 -u"De
v" --ul-style=filled --ul-font=Times-Roman120 -P top2
```

Fig. 2 – DRIVER.TOP2

```
execute "SH -c ' ": '/usr/local/bin/printout.converter " ':JOB.ID:' " " ':SRC.DIR:' " " ':DEST.
DIR:' " " ':PG.ORIENT:' " " ':REMOVE:' " ":" ' "
  execute "SH -c ' ": '/usr/local/bin/mail_attach.pl " ':EMAIL.RECIP:' " " ':SUBJECT:' "
"ATTACH=':SRC.PATH:' " " ':REMOVE:' " ':" ' "
```

Figure 3 - mvBASIC

cial knowledge on the end users' part, a real consideration when the report could be going to any one of 15,000 employees; from the President to the warehouse workers.

This didn't solve the email problem but did have a hidden benefit: when emailed, the pdf report is much smaller than the original text report. I have used the same process to create the driver for my Unix/Linux printers that print UniVerse reports *[Figure 2]*. Most printers on the network are now Windows printers that the Unix/Linux world borrow. They are postscript enabled.

This also allows me to put a watermark on any printouts coming from test and development systems. That solved a worry of mine about test/development reports escaping into the "wild" with no indication of where they originated.

So, now I have the ability to get text files printed, but the email was still a problem. I was using Google to search for some way to solve the email problem and found a Perl program that packages attachments and emails them. The program is called `mail_attach.pl` and was written by Steve Ho. (`mail_attach.pl` can be downloaded from https://www.intl-spectrum.com/resource/786/default.aspx)

It handles all of the header elements for the developer. If you are sending attachments, it handles that, too. The developer only needs to know what parameters to send to mail_attach.pl. I have made a few enhancements to it. We also added code to "jail" emails from our test and development systems. There is a configuration file for `mail_attach.pl` and we added code and parameter so we could redirect all emails to a single mailbox, retaining the original address.

I've included two example commands from the mvBASIC programs *[Figure 3]*.

Combining these techniques gave us a single, standardized way to send e-mails automatically. **IS**
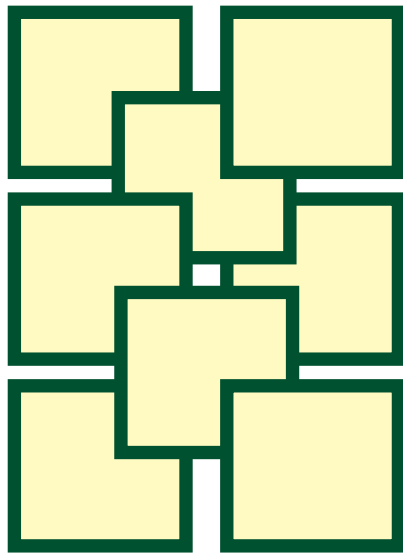
OpenQM Magic!

BY GRAEME MCGILLIVRAY

TOC, TOC, TOC

The Theory of Constraints (TOC) has, as a main tenet, the idea that every system is limited by its constraints. TOC has a methodology for finding those constraints and overcoming their limitations. While I'm not sure that the developers of OpenQM have used the Theory of Constraints, I am sure that they have identified certain limits in the MultiValue programming model. The technology that has been introduced into OpenQM that has got me talking about TOC is objects.

But before I explain the potential of this marrying OO (Object Oriented) to MultiValue, there is one other issue I want to raise relevant to TOC: The classification of constraints. The most obvious constraints in any system are physical ones, but the most important *and* significant constraints are the non-physical limitations (policy constraints). Vendors and users of any data model, including SQL, have a major policy constraint: Resistance to change. I'm impressed that OpenQM has overcome this resistance by implementing important new features.

> *...developers do things like "borrow" a little-used (from their perception) common variable as a quick fix for an extra data transmission between functions.*

What's so Important About Objects Anyway?

One of the best theorists on this subject is Bertrand Meyers. To quote from his book "Object-Oriented Software Construction" he states that "the main problem (in software) is the data transmission structure of the system." He then continues "the background to object technology, as presented at the beginning of this book, is a battle between the function and data (object) aspects of software for control of the architecture. In non O-O approaches, the functions rule unopposed over the data; but then the data take their revenge."

I'm sure you have all seen applications with large commons, which usually exist because of the various and complex data transmissions between functions. The easiest way to manage this is to dump all the various variables into one large common. Then, unfortunately, developers do things like "borrow" a little-used (from their perception) common variable as a quick fix for an extra data transmission between functions. I'm sure you have all suffered from this side-effect when debugging a function and being frustrated by finding the data you didn't expect in a common variable.

So Back to the Question "Why Objects?"

Law of Inversion
If your routines exchange too many data, put your routines in your data.

Instead of building modules around operations and distributing the data structures between the resulting routines with all the unpleasant consequences that can occur, object-oriented design does the reverse. It uses the most important data-types as the basis for modularization, attaching each routine to the data-type to which it

relates most closely. The law of inversion is the key to obtaining an object-oriented design from a classical functional (procedural) decomposition. Analyzing data transmissions is a good way to detect and correct design flaws.

Now For An Example

This example comes from Oracle for Objects. Yes, it's an SQL product, but Oracle has moved on from the traditional SQL implementation, and introduced changes to make their database conform more closely to E.F. Codd's original theoretical model. But before presenting the problem and the OpenQM solution, we need one more bit of background information to show why objects are important to the multivalued model.

C.J. Date and Hugh Darwin in their book "Foundation for Object/Relational Databases, The Third Manifesto" present the argument why the SQL vendors have made a complete mess of implementing the relational model and what the model actually represents. A key tenet is that "Databases of the future will need to contain much more sophisticated kinds of data than current commercial ones typically do." The second tenet is that "a relation (the model terminology as distinct from the term table which is implementation terminology) has two parts, a heading and a body, where the heading is a set of column-name: domain-name pairs, and a body is a set of rows that conforms to the heading."

The crux of the issue is in how we define the domains. In SQL its simple, single-valued data types such as integer, varchar, date etc. In the MultiValue model it is numeric or text arrays representing the domain values. In Oracle Objects we can have the construct shown in Figure 1.

Note that Date and Darwin state that a "domain is the equivalent to a class." And yes, the standard MultiValue model could handle this, something like Figure 2.

A record structure such as

```
L1]Auckland^p1\Bill\1234|9876]P2\Jack\654|9977^
```

Updating the second phone number of the second person would require something like:

```
Read record from fvar_Office,"O001"then
    Record<2,2,3,2> = {new value}
end
```

A practical solution, but certainly not intuitive. And what happens if person two is an occupant of more than one office, can we be sure that they both get updated – to name just one issue.

Before converting this into OpenQM, an explanation of the Oracle Object steps above is necessary:

1. *Create type* is the definition of a Domain. This type definition is an array definition, the equivalent of stating DIM PHONE_LIST_ARR(10)

2. This type definition is of a person Domain (or class), the equivalent would be a dictionary, something like**:**
@ID = idno
   <1> = name
   <2> = phone(10)

3. This type definition is for a location Domain (or class), the equivalent would be a dictionary, something like:
@ID = building_no
<1> = city

```
1.  CREATE TYPE phone_list_arr AS VARRAY(10) OF VARCHAR2(20);

2.  CREATE TYPE person_typ AS OBJECT (
        idno            NUMBER,
        name            VARCHAR2(30),
        phone           phone_list_arr);

3.  CREATE TYPE location_typ AS OBJECT (
        building_no   NUMBER,
        city          VARCHAR2(40) );

4.  CREATE TYPE nt_persons_typ AS TABLE OF person_typ;

5.  CREATE TYPE office_typ AS OBJECT (
        office_id    VARCHAR(10),
        office_loc   location_typ,
        occupant     nt_persons_typ );

6.  CREATE TABLE office_tab OF office_typ (
        Office_id       PRIMARY KEY );
```

Fig. 1

Office-ID <@ID>	Office_loc <1>		Occupant <2>		
O001	Building_no	City	Idno	Name	Phone
	L1	Auckland	P1	Bill	-1234 -9876
			P2	Jack	-654 -9977

Fig. 2

4. This type definition is for a *nested table* (think of the intersection of a row and column, not just holding a single value, but a table of values). The nearest equivalent in the multi-value model would be equivalent to:

<2> = nt_person_type(x,y) – a two dimensional array in a file attribute. Another equivalent would be rows = VM's, and columns in the row = SVM's.

5. This type defines a domain (a class), the equivalent would be a dictionary, something like:

@ID = office_id
<1> = office_loc (of type location_typ)
<2> = occupant(x,y) (of type person)

6. This table is the domain type of office_typ

The OpenQM Solution

The Person class is definition can be found at: https://www.intl-spectrum.com/resource/782/default.aspx.

The location class definition can be found at: https://www.intl-spectrum.com/resource/782/default.aspx.

You'll notice that both these classes have data persistence by having their own files to store the object contents.

These classes satisfy the Law of Inversion in that these file I/O routines are attached to the data-type to which it relates most closely. This is a major change from having one monolithic program carrying out a multitude of file I/O for a range of files.

Both classes have a INIT routine to read and initialise the object contents. Both classes also have a X = ME->SAVE call to the SAVE routine whenever any internal value changes.

The office class definition can be found at: https://www.intl-spectrum.com/resource/782/default.aspx.

Class Office defines the domain Office_Typ.

In the property definition of the class are two private object pointers,

the first for a single location object and the second a dimensioned array for the multiple people objects. The structure of the People Array is

```
Person ID 1        Person 1 Object Pointer
Person ID 2        Person 2 Object Pointer
```

1. The private location object pointer in the Office class holds the reference to the created Location Object.

2. OpenQM allows a dimension array to be re-dimensioned on the fly, so it can grow to accommodate the correct number of child person object pointers.

3. If the office record does not hold a location, then create a location object using the Location object pointer, record its details in the Office object and ask the office object to save itself because its content has changes. Note: The routine does NOT specifically force the Location object to be saved. That is implicit within the Location object, it is its responsibility to save itself if its content changes.

4. To change a property within the Location object, the Office object just sends a message to it {X = LOCATION.OBJ->SET.MEMBEROF(IDNO)}. It is up to the Location object to manage its own internal state, which it does.

Finally, Let's Put This All Together

In OpenQM, the syntax to build a file using the Office type is not yet possible, but if this was controlled by a program, then the structure would be something like Figure 3.

We now have a collection container to hold all the office type objects (this is the equivalent of the Oracle Objects statement.

```
OPEN "OFFICE.TYPE" TO FVAR ELSE
    ERR = "UNABLE TO OPEN OFFICE.TYPE FILE"
    STOP ERR
END
*
DIM OFFICES(1,2)
MAT OFFICES = ""
* - select the OFFICE.TYPE file and load up column 1 of
OFFICES(row,col)
HUSH ON
EXECUTE "SELECT OFFICE.TYPE BY @ID" RTNLIST LISTVAR
HUSH OFF
MAX.KT = 0
LOOP
    READNEXT ID FROM LISTVAR ELSE ID = ""
WHILE ID DO
    MAX.KT += 1
    * dynamically re-dimension dimension array to hold object pointers
    DIM OFFICES(MAX.KT,2)
    OFFICES(MAX.KT,1) = ID
REPEAT
```

Fig. 3

```
6. CREATE TABLE office_
tab OF office_typ (
    office_id
PRIMARY KEY );
```

Now perhaps the user wishes to update a specific Office record, something like the following as the example in Figure 4.

```
DISPLAY "enter office number ": ; INPUT OFFICE.ID
IF OFFICE.ID = "" THEN STOP
VA = "" ; FOR I = 1 TO MAX.KT ; VA<I> = OFFICES(I,1) ; NEXT I
LOCATE(OFFICE.ID, VA; KT ; "AR") THEN
    * check whether object created
    IF OFFICES(KT,2) = "" THEN
        OFFICES(KT,2) = OBJECT("CLASS.OFFICE.QM")
        OFFICES(KT,2)->INIT(OFFICE.ID)    END
    {… add logic here …}
END
```

Fig. 4

How do we add logic once the office object is created?

To read the Office's Location's City:

```
D = OFFICES(KT,2)->GET.LOCATION.CITY
```

To update the Office's Location's City:

```
OFFICES(KT,2)->SET.LOCATION.CITY("PAPATOETOE")
```

Obviously there is much more that can be done than what I have written about here. Let me summarize the key points covered.

• Overcome your own policy constraints and investigate changes which can benefit you.

• When a design or an existing application is becoming too complicated, use the Law of Inversion and analyze the data transmissions between functions.

• Where possible, isolate the most important data types and by using the object features of OpenQM, attach the appropriate routines to the data type.

• A column's domain does not have to be a simple, single value as per the SQL model, or an un-described array as in the traditional MultiValue model.

• Complex domains can de-defined as classes, where the class has the responsibility to manage its own affairs.

• Using classes to define domains overcomes the problem of poor clarity. What does `record<2,2,3,2>` mean?

• When using classes, make them responsible for their own actions. Goggle "Design by Contract" (a Bertrand Meyer concept) which formalizes the contract between a class and its clients.

• You can simplify the use of class domains by using a construct such as:

```
Class Simple
    Public WorkArray1(100,100)
    Public WorkArray2(50,50)
    Public Var1, Var2, Var3
    End
```

And the use of such a simple class would be:

```
Obj = create.object("Simple")
Obj->WorkArray(3,3) = value1
X = Obj->Var3
Obj = ""  ;* destroy object pointer when
finished
```

• And of course, download OpenQM and take the object extension for a ride and enjoy. **IS**

GRAEME MCGILLIVRAY is an independent consultant in Auckland, New Zealand. In a former life, in the Pick glory days in NZ, he was the president of the NZ Pick Users Group.
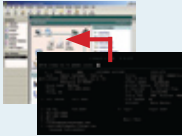
MARKETPLACE

ACCOUNTING

Natec Systems
www.natecsystems.com | nrector@natecsystems.com

COMPLIANCE

SJ+ Systems Associates
www.sjplus.com | sjoslyn@sjplus.com

CONSULTING

Drexel Management Service
www.drexelmgt.com | dconboy@drexelmgt.com

Execu-Sys, LTD
www.eslny.com | mh@eslny.com

HDWP
www.HDWP.com | results@HDWP.com

Modern MultiValue, LLC
www.ModernMultiValue.com | info@ModernMultiValue.com

PICK Programmers Shop
www.pickprogram.com | brian@pickprogram.com

Precision Solutions
www.precisonline.com | Kevin@PrecisOnline.com

DATABASE

Ladybridge Systems Ltd
www.ladybridge.com | sales@Ladybridge.com

REPORTING

Brian Leach Consulting, LTD
www.brianleach.co.uk | brian@brianleach.co.uk

TERMINAL EMULATOR

Zumasys
http://www.zumasys.com/products/accuterm/

LETTERS TO THE EDITOR

Have an opinion on an article: Agree, disagree, or enhancement to an article from a previous issue? International Spectrum and our authors are interested in hearing from you!

E-mail: editor@intl-spectrum.com

WANT TO SEE A SPECIFIC TOPIC?

International Spectrum is looking for writers, feedback, and topic ideas. We all have specific topics and issues that we need answers to find solutions for. Send us an E-mail with topics you would like to have covered in the magazine or on the website.

E-mail: nathan@intl-spectrum.com

WANT TO WRITE?

Expand your professional credentials, and provide us with an article.

Give us a rough and ugly outline, and we will help you refine it, proof it, and make it press ready. Or you can give us something polished, proofed, and press ready to publish.

Share your thoughts and expertise with over 10,000 fellow MultiValue developers and users.

E-mail: editor@intl-spectrum.com

NEED A MENTOR?

Mentors give developers the ability to ask industry experts for direction, code examples, and/or just ask them to see if something makes sense. Sometimes, all you need is a resource or example to start or complete a project.

Check with us to see who is available for mentoring, and how you can take advantage of it to save your business or company money.

E-mail: nathan@intl-spectrum.com

WANT TO BE A MENTOR?

We have many retired or semi-retired professionals out there that would love to share their knowledge of MultiValue development. If you are one of them, please contact us to see what mentoring is all about.

E-mail: nathan@intl-spectrum.com

IS.HASH.MD5

Generating MD5 Values with Generic MultiValue BASIC

The MD5 message-digest algorithm is a widely used cryptographic hash function producing a 128-bit (16-byte) hash value, typically expressed in text format as a 32-digit hexadecimal number. MD5 has been utilized in a wide variety of cryptographic applications and is also commonly used to verify data integrity.

While not as secure as SHA1, it is still used in many places for data integrity, version control, and other features that need unique one-way signatures.

MD5 Function

MD5 processes a variable-length message into a fixed-length output of 128 bits. The input message is broken up into chunks of 512-bit blocks (sixteen 32-bit words); the message is padded so that its length is divisible by 512. The padding works as follows: first a single bit, 1, is appended to the end of the message. This is followed by as many zeros as are required to bring the length of the message up to 64 bits fewer than a multiple of 512. The remaining bits are filled up with 64 bits representing the length of the original message, modulo 264.

```
MD5("The quick brown fox jumps over the lazy
dog")
9e107d9d372bb6826bd81d3542a419d6
```

MD5 will detect even small changes in the string and cause the returned hash value to change. You can always read more about MD5 at:

https://en.wikipedia.org/wiki/MD5

MultiValue BASIC Hashing

Most versions of MultiValue BASIC do not include built-in hashing functions. This requires developers to create their own. We can use Pure/Generic MultiValue BASIC, but you may see a performance hit depending on how bit math is handled. For hash values like MD5, the performance issues of generic bit math does not seem to be a problem.

Dave Meagher has done all the work to provide hash functions like SHA1 and MD5 in his FOSS4MV/mvCrypt code on BitBucket:

https://bitbucket.org/foss4mv/mvcrypt

The IS.HASH.MD5 program attached to this article used Dave's mvCrypt MD5 code, but formats it so that it is interchangeable with other versions of the IS.HASH.MD5 code on the International Spectrum Website.

Example

```
MD5.MESSAGE = "The quick brown fox jumps over
the lazy dog"
CALL IS.HASH.MD5(SHA1.MESSAGE,HASH.VALUE)
*
TEST.VALUE = "9e107d9d372bb6826bd81d3542a419d6"
CRT HASH.VALUE : " =" : TEST.VALUE :" - "
IF (HASH.VALUE EQ TEST.VALUE) THEN
   CRT "Ok"
END ELSE
   CRT "Failed"
END
```

Output

```
9e107d9d372bb6826bd81d3542a419d6 =
9e107d9d372bb6826bd81d3542a419d6 - Ok IS
```