

A response to R. Camps' article
"Domains, Relations and Religious Wars"

by

C. J. Date

Voice: (1) 707/433-6523

Fax: (1) 707/433-7322

September 15th, 1996

Since it quotes extensively from writings of my own, I feel obliged to respond to the article "Domains, Relations and Religious Wars," by R. Camps (*SIGMOD Record* 25, No. 3, September 1996). In that article, Camps is clearly suggesting (among other things) that my definition of the term "domain" has changed over the years. I agree, it has! But Camps goes on to say:

"... considering that [Date's book *An Introduction to Database Systems*] was the *bible* [Camps' italics] where most university graduates all over the world learnt, I believe that Date can be held partly responsible for the lack of implementation of domains [in today's SQL DBMSs]."

I reject this accusation, strongly; in fact, I object to it on several levels at once.

- First of all, the lack of domain support in today's SQL DBMSs is surely due, first and foremost, to the lack of domain support in SQL itself -- and I can say with some assurance (and some feeling) that my influence on the design of SQL has never been more than minimal.
- Second, the question of which books happen to be selected by universities as "bibles" is quite beyond my control (I might wish it otherwise, but it isn't).
- Third, while it's true that I've made some changes to my definition of "domain" over the years, I would argue that those changes reflect merely improvements in my own understanding, not a rejection of what I once believed. In other words, the changes have, I think, always been "backward compatible"; a DBMS that implemented domains as I first described them would in no way be precluded from supporting domains as I see them now (i.e., as data types, in the full sense of that term).
- Fourth (as Camps does admit in his article, but too late), those definitional changes have all been made very much in the spirit of Bertrand Russell, who wrote:

"I am not myself in any way ashamed of [having changed my

opinions] ... The kind of philosophy that I value and have endeavoured to pursue is scientific, in the sense that there is some definite knowledge to be obtained and that new discoveries can make the admission of former error inevitable to any candid mind. For what I have said, ... I do not claim the kind of truth which theologians claim for their creeds. I claim only, at best, that the opinion expressed was a sensible one to hold at the time ... I should be much surprised if subsequent research did not show that it needed to be modified."

As I said when I first quoted these remarks (in the 6th edition of *An Introduction to Database Systems*): "Readers ... will find that I too have changed my opinions on many matters (and will no doubt continue to do so). I hope they will accept the remarks quoted above as adequate justification for this state of affairs. I share Bertrand Russell's perception of what the field of scientific inquiry is all about, but he expresses that perception far more eloquently than I could hope to do."

- Fifth, I object strenuously to the suggestion -- reflected not only in the extract from Camps' article quoted above but also in the title, in one of the two epigraphs, and indeed throughout the article -- that the debate over data models is a "religious" matter. (Though in fairness I have to say that Camps is not alone in characterizing the debate in such a way.) Ever since its inception, database management, like the rest of computing, has aspired to be a scientific discipline. In this connection, the relational model represents a great leap forward, being founded as it is on logic, not religion. Why then do database professionals (some of them, anyway) so often talk as if logic and science had nothing to do with the matter? Do they perhaps not wish their discipline to be seen as scientific? If not, why not? (In this connection, note Bertrand Russell's remarks, already quoted, on "the kind of truth which theologians claim for their creeds.")

In addition to the foregoing, there are some specific points in Camps' article that I feel need some further response:

- "[First] normal form ... was one of the most surprising, and most difficult to accept, of the RM principles ... [it was] viewed ... as a strong restriction and a step backwards ... [and] was a major battlefield in the confrontation between DBTG-CODASYL and RM supporters."

Frankly, I don't remember much of a fight over this one. *Au contraire*, in fact: I can remember the then chair of DBTG, Tax Metaxides, saying at a conference that he thought first (and indeed third) normal form was a good idea. No, the big argument was over *manual vs. automatic navigation* ("automatic navigation will never perform," etc.).

- "[During the] bitter fight between the relational camp and the

DBTG-CODASYL camp ... [the] relational camp was mainly composed by IBM ... Under the cover of technical and scientific arguments, relevant commercial interests were being disputed."

It's possible that "the DBTG-CODASYL camp" might have had some commercial interests at heart (and at stake) during this "bitter fight," but it's nonsense to suggest that "the relational camp" (or IBM) did so. There weren't any relational products! (certainly not IBM products) -- nor were there even any such products under development at the time. (The "Great Debate" between the two camps was held in 1974. IBM's first mainstream relational product, SQL/DS, was announced over seven years later, in 1981, and DB2 wasn't announced until 1983.)

- "Would it not be more reasonable to look for the cause [of lack of commercial domain support] in the RM's lack of a type system?"

No, it wouldn't. The relational model very deliberately does not prescribe any specific data types. As Hugh Darwen and I wrote in a recent article, "Introducing ... *The Third Manifesto*" (*Database Programming & Design* 8, No. 1, January 1995), the question as to what data types are supported is orthogonal to the question of support for the relational model. Or, more catchily, types are orthogonal to tables. Thus, the relational model does require support for domains, which means (by definition) that users should be able to define their own; it doesn't spell out exactly which ones they must define.

- "How can Date claim that [user-defined types are] not an extension to the RM?"

But they're not! See the previous bullet item.

- "[Had] the religious war of the late Seventies between DBTG and relational camps been a rational debate, free of fanaticism and commercial interests, today's scenario would be quite different."

It was not a "religious" war. And, speaking here as a member of "the relational camp," I would say that that camp, at least, did do its best to make the debate as "rational" and as "free of fanaticism and commercial interests" as possible.

In closing, let me just say this: What's really important, of course, is not **who's** right or wrong, it's **what's** right or wrong. I concur with Camps that "it is always profitable to review our old opinions and beliefs," but I think it's more important to agree on how best to move forward. The work of Hugh Darwen and myself on *The Third Manifesto* is, and continues to be, aimed at this goal.