DRIFT/VISC: DISTRIBUTED, REDUNDANT, AND INHERENTLY FAULT TOLERANT MICROPROBE CONCEPT

MARK ARNOLD*

**Department of Computer Science and Engineering, Lehigh University, USA*
maab@lehigh.edu

ED CHESTER

Aerospace Research and Technology Centre, Barcelona, Spain
ed.chester@ctae.org

ABSTRACT

MEMS (Micro-Electro-Mechanical Systems) are already a key enabling technology for a number of technical challenges in the space sector, delivering mass-specific and power-specific performance and capability that exceeds conventional technologies by orders of magnitude. Since MEMS reuse IC fabrication technology, they can be readily mass produced. System-on-Chip has already been considered as a paradigm leading to single-board spacecraft (*e.g.* [1]), and spacecraft-on-a-chip [2]. In this paper, we assume this interest and trend will continue. This paper proposes a possible way of using multiple probes to form a more capable observation and processing system, and we consider a highly distributed computer composed of distinct probes as a mission design concept, but from the atypical standpoint of computing science rather than spacecraft engineering.

We eliminate the tangible presence of the memory from the probe, and only fabricate the hardware for each probe CPU on the probe together with the radio circuits already needed to relay information, thus taking the concept of RISC (Reduced Instruction Set Computer) a step further to form VISC – the Vacuum Instruction Set Computer. The principal novelty here is that the 'storage' for the VISC is the RF signal as it propagates. During the radio propagation time between probes, many distinct bits may be in transit. While any particular bit may be lost or corrupted, error correcting codes allow such errors to be detected and corrected, using minimal extra circuitry. The deployment of a large number of such probes (for example, into an atmosphere) provides a longer time-scale observation of structure and dynamics than can be achieved with individual complex probes. These relatively inexpensive devices also form a distributed, redundant and inherently fault tolerant (DRIFT), processing platform for handling the collected data.

A concept design for DRIFT is presented, with a functional model of VISC and its instruction set, together with simulations of probe interactions and system level performance estimates. Engineering challenges and the applicability to exploration is considered.

Nomenclature

a	Bandwidth fraction allocated to memory
D	Mission range [m]
d	Probe separation [m]
$d_{i,j}$	Data bit j of word i
$e_{i,j}$	ECC bit j of word i
f	Bit rate of interprobe channel [s^{-1}]
h	Altitude [m]
i	Memory address
γ	Lorentz factor
m	Mass [kg]
n_a	Address width [bits]
n_d	Data width [bits]
n_h	Communication link hops
n_p	Number of VISC nodes
n_r	Total wordlength [bits]
n_t	DRIFT train length
n_v	Memory cycles
P_S	Static and payload power dissipation [W]
P_L	Logic power dissipation [W]
P_T	Transmission power [W]
R	Planet / Target radius [m]
r	Redundancy factor
t	Memory latency [s]
T	Duration [s]
V	Relative radial velocity [ms^{-1}]
X_L	Logic power per bit [W]
X_T	Transmission power per bit [W]

1 INTRODUCTION

Space is often associated with some of the grandest structures of human endeavor. 'Big' is an understatement characterizing the nature of space, and also space programs. Here, we consider the new paradigm of technological achievement — the very small; elegant compactness; extreme integration; a generally broad technological shedding of mass and power consumption, of material and energy.

A recent achievement of microtechnology has been the creation of Micro-Electro-Mechanical Systems (MEMS), integrating mechanical elements, sensors, actuators, and electronics on a common silicon substrate through microfabrication technology. MEMS thus allow common spacecraft functions, such as attitude control and propulsion, to be integrated on a microchip [3].

Despite these efforts, it has proven difficult to design and qualify MEMS components fast enough to include them in space missions [4,5]. TRL raising and qualification is an expensive and complex undertaking, long term reliability and performance are yet to be demonstrated, and susceptibility to radiation-induced upsets, are all among the factors limiting the uptake of MEMs as an enabling technology for space systems.

The concept discussed in the remainder of this paper addresses some remaining limits. Within even the simplest spacecraft, the necessary functions of storing and processing data, and of communicating with a remote receiver, are implemented using the same basic methods as the largest vehicles and systems. Table 1 illustrates the specifications of computers used in some historic spacecraft, with the intention of highlighting that extremely simple machines have successfully executed complex tasks in remote locations. Clearly more advanced computers are available to spacecraft designers today, but our objective is simplicity rather than raw performance.

Spacecraft/CPU	*Power / MIPS*	*Wordsize*	*Storage /Kwords*
Apollo/AGC	0.1	16	5
Galileo/ATAC-16, 1802	0.5	16	64
Sojourner/80c85	0.1	8	8
STS/GPC	0.2	36	32
MILSTAR/TDY750	0.45	16	512

Table 1: Historic Spacecraft Computers

Conventional computers have been adapted for use in harsh environments with high reliability, incorporating features of fault tolerance and error checking and correction (ECC). Onboard data storage has evolved from magnetic core memory, via tape and hard disks, to solid state mass memory in the same way as terrestrial computers, but lagging at least one performance generation behind the state of the art due to qualification delay.

2 A PLATFORM OF DISTRIBUTED MULTIPROBES

Distributed, independent probe survey is used in terrestrial applications, and has been proposed for future solar system mission concepts. Drop probes carrying locators can provide wind, atmospheric structure and dynamics, and even chemical composition

information. Conventional descent probes, such as Vega [6] or Huygens [7], have returned large volumes of valuable science data from complex instruments, but addressing limited target regions.

Conversely, very simple probes with limited capabilities can be distributed to obtain boundary values for global measurements. Designs for small distributed probes are already well advanced (*e.g.* for Venus [8]), which conceptually fall between conventional descent probes and highly distributed microprobes in terms of their tradeoff design point in science value, technology readiness, deployment numbers, and probe size.

Through evaluating the limiting case of how far the size/complexity trade can be pushed, we note that aside from payload to measure some observable parameter, an increasingly dominant portion of a microprobe design is allocated to support functions – data handling, processing, and communications. To mitigate this and facilitate further miniaturization, we divide the processor and communication functions across a highly distributed set of microprobes. Each of a great many microprobes acts as a node in a digital circuit.

To investigate the viability of this idea, a computing science approach is adopted, before returning to engineering considerations, thence a futuristic extension. The following section describes the design of a data handling and processing system (actually a simple computer) across this platform of a large number of distributed nodes.

3 VACUUM INSTRUCTION SET COMPUTING

Conventional computers have two main components: the Central Processing Unit (CPU), which performs calculations and makes decisions, and the memory, which holds a program that instructs the CPU what to do, and holds the results of calculations and other data. Although some commercial CPUs are highly complex with many instruction types, Reduced Instruction Set Computers (RISCs), popularized during the last two decades, use much simpler CPUs to accomplish the same tasks with somewhat longer programs consisting of fewer instruction types.

Having reduced the hardware complexity of the RISC CPU, the memory becomes the single largest component of the computer. In this paper, we take the idea of a RISC one step further, which we refer to as a Vacuum Instruction Set Computer (VISC). We eliminate the tangible presence of the memory from each microprobe, and only fabricate the hardware for the CPU on the probe together with the necessary radio circuits needed to relay information through a network of such devices.

The unusual innovation here is that the 'memory' for the VISC concept is *only* propagating radio signals. The VISC uses a small number of microprobes to relay the contents of memory between each other. This is analogous to a digital shift register. Obtaining any particular bit may require waiting for all of the other bits to pass by first. A *serial* computer is not a new idea — some of the earliest computers used such delay-line memories [9].

The difference here is that the bits are actually being 'stored' in space. With suitable separation, it takes a very small fraction of a second for a radio signal to propagate from one probe to another. During that time a large number of distinct bits may be in transit in each direction. As normally experienced on space communication links, any particular bit may be corrupted, but Error Correcting Codes (ECC) allow such errors to be detected and corrected, and routine techniques such as those captured by the CCSDS standards [10] reduce practical bit error rates (BER) below acceptable minimum values for useful communications (BER of 10^{-7} is not unusual). ECC implies an overhead of extra circuitry and some part of the effective link bandwidth, but it maintains overall system performance in the presence of an imperfect communication channel.

Temporarily ignoring physical implementation or deployment, we assume that a very large number ($>10^6$) of microprobes forms a single DRIFT – a distributed, redundant, fault tolerant VISC

deployment. The machine comprises a distributed CPU and serial memory held in radio waves; together they form a very large state machine with 2 basic functional architectures as follows.

Fixed Function Architecture

This paper will assume that there is a static mapping of each node to a particular role in the DRIFT network, *e.g.* CPUs, repeaters, *etc.* The basic scenario imagined is that each node is a probe moving relative to a target of observation, and the target is assumed to be a planet, moon, or other small body. Each node retains its unique function for as long as it remains in the network, and this function is lost when the node leaves the network (through separation, atmospheric entry, failure, *etc.*).

Fixed Region Architecture

Since the DRIFT swarm could be moving relative to the target, an alternative architecture could allocate a region of space to a particular function, with a through-flux of probes that implement the desired function in an approximately stationary region relative to the target. To place this option into context, solar-powered microprobes (*e.g.* [11]) may maintain the active VISC on the day-side of a target to allow continuous operation even after parts of the network have passed out of sunlight, or entered an atmosphere, or left communications range, *etc.*

The state of the VISC would stay within one group of probes for a time proportional to their velocity. After that time, transferring the state from one group of probes to another group (now roughly located where the first group was earlier) occurs naturally with the serial memory we propose. This alternative is not considered further in this paper, but presents some interesting possibilities for future development of highly redundant probes.

The serial memory makes each VISC slow by any normal standards of computation. This is compensated for by the possibility of executing a very large number of VISC threads in parallel, providing sufficient computational power to perform complex tasks such as image compression. Multi-threaded VISCs allow for intelligent compression of data (for example, measurements gathered from the target object). The assembly, processing and transmission of the data requires only a few digital logic gates in each node.

Although the power of parallel processing was known to bit-serial pioneers a half-century ago, the difficulty of parallel software prevented implementation of the concepts we propose. In the 1980s, Hillis [12] used a large quantity of bit-serial processors (with semiconductor RAM rather than the delay-line memory we propose) for his Connection Machine, but its programming paradigm was not widely accepted. Only in recent years have parallel-software paradigms like OpenMP [13] and CUDA [14] reduced the difficulty of programming a multi-threaded architecture akin to VISC.

The research questions posed are then straightforward: what is a minimal VISC implementation, what circuitry is required in each microprobe, which instructions does it support, and how can this concept be useful? These are addressed in the following sections.

4 VISC SIMULATOR

To estimate the smallest logic that could implement the VISC concept, we modified a Verilog model for a minimal bit-serial CPU (the PDP-8/S) [15], and implemented only the most essential instructions for general-purpose capability (*i.e.* add, logical AND, store, complement, test and branch). We extended this model by implementing simple error-correcting codes, as would be essential in actual deployment. To model the delay-based memory easily, we used a feature of the Verilog simulator (non-blocking assignment with delay) that allows a large number of simulated bits to be in transit between the separate Verilog models of ‘CPU probe’ and a ‘repeater probe’.

To verify the correctness of the design, division-by-repeated-subtraction has been implemented in software [15], and applied within the VISC by initializing this task from an 'oracle' craft that broadcasts the initial state of the VISC (analogous to uplinking a software patch from an Earth ground station).

From that point on, the oracle remains silent, and the feedback between the VISC and repeater maintains the memory. To validate the ECC implementation, some bits of the initial bitstream provided by the oracle are intentionally corrupted. The Verilog model of the controller is able to correct these bits, and execute the program successfully. This choice of task ensures that all the functionality of the VISC Verilog model is exercised in simulation.

To estimate the required silicon area for a real implementation in a contemporary VLSI technology, we used the VITO tool that translates so-called 'implicit Verilog' into a *one-hot* controller [15,16]. Figure 1 shows a layout generated for the controller, which has the active elements listed in Table 2:

Component	*Symbol*	*Quantity*
OR	X	22
DEMUX	d	20
FLIPFLOP	f	25

Table 2: VITO Synthesis Summary of VISC CPU

The parts of in boldface implement the error-correction logic, which constitutes about 50% of the 21×23 circuit layout shown. Commercial synthesis would provide a more compact layout. Design hardening techniques, in particular the addition of high-dopant guard rings around each transistor to protect against latch-up, will have significant penalty in layout area [17].

```
/>X<~<1
^ |
| f
^ |
| X<-<-<-<-<-<-<->->->-<-<-<-<-<-<-<-<-<\
^ |                                      ^
| d>->->->->->->->->->->->->->->\         |
^ |                            |         ^
| d>\                          f         |
^ | |                          |         ^
| | f                          X<-<-<-<-<\ |
^ | |                          |         ^ ^
| | X<-<-<-<-<-<-<->->->-<\ d>->->\   | |
^ | |                      | |     |   ^ ^
| | d>->\                  ^ f     f   | |
^ | |   |                  | |     |   ^ ^
| | f   f                  ^ d>\   X<f | |
^ | |   |                  | | |   | ^ ^ ^
| X</   X<-<-<-<-<-<-<-<\ ^ | f   d>/ | |
^ |     |               | | | |   |   ^ ^
| d>\   d>->->->\       ^ ^ | X<f |   | |
^ | |   |       |       | | | | ^ |   ^ ^
| | f   d>\     f       ^ ^ | d>/ d>\ | |
^ | |   | |     |       | | | |   | | ^ ^
| X</   | f     X<-<f   ^ ^ | f   X</ | |
^ |     | |     |   ^   | | | |   |   ^ ^
| d>\   | x     d->/    ^ ^ X</   \>-> / |
^ | |   | |     |       | | |           ^
| | X<f | d>\   d>\     ^ ^ \>->->->->->/
^ | | ^ | | |   | |     | |
| | d>/ X</ f   | d>\   ^ ^
^ | |   |   |   | | |   | |
| X</   d>\ X<f | | f   ^ ^
^ |     | | | ^ | | |   | |
| d>\   | f d>/ | | f   ^ ^
^ | |   | | |   | | |   | |
| | f   | x f   | X</   ^ ^
^ | |   | |     | |     | |
\<X</   | d>\   X</     ^ ^
        | | |   \>->->->/ |
        | f f             ^
        | |               |
        X</               ^
        |                 |
        \>->->->->->->->->/
```

Figure 1: Controller Synthesis using VITO

5 ENGINEERING CONSIDERATIONS

Assuming that the substrate thickness and dimensions for each required gate are 1μm (very conservative compared to current fabrication techniques at <100 nm), the volume required for implementation of the controller of Figure 1 would be $4.8 \cdot 10^{-10}$ cm^3. The device would need to include other components (most significantly, data and address registers not shown in Figure 1). Even with these, the mass of the digital logic for a minimal VISC would be of the order of a few nanograms.

For communications based upon ordered resonating nanotubes [18], an additional comparable mass is expected. The mass for the power supply will depend heavily upon the type used. Solar powered probes [11] will probably require a much larger area

devoted to power supply than to digital logic.

To model the storage capacity of the space between two probes, whose inter-probe distance is d, we assume their relative velocity is zero. We also assume that there are at least two separate communication channels that allow full-duplex communication. The number of hops n_h for the most straightforward VISC implementation is 2, so that at any moment half of the bits are travelling to the first probe (which has the hardware for the VISC) and half of the bits are travelling to the other probe (which is a repeater). If $n_h > 2$, $(n_h - 1)$ of the probes could be repeaters arranged in a helical structure allowing the first and last repeater to be at distance d from the VISC CPU probe. For sufficiently large n_h, a single VISC could form a closed network around a planet, albeit with very long latency.

The communication channels allow transmission of f bits per second. Assuming negligible repeater delay, the total delay for a particular bit to have travelled through the repeater(s) and back is given by $t = n_h d/c$. The bitstream has a total throughput capacity of $f{\cdot}t$ bits; however, we assume that only an allocated fraction, a, of this raw stream is budgeted for use as a memory bank, with the remaining (1 - a) used: *i*) to cope with interprobe distance, clock and power-supply variations, and *ii*) to allow long distance, navigation, and housekeeping communication within the DRIFT network.

The allocated bank needs r times as many bits as an ideal memory to provide redundancy for error correction. The VISC has a conventional computer architecture, with n_a address bits, and n_d data bits embedded with redundancy within n_r bits. The ratio of redundancy, $r = n_r / n_d$, is unity if no error correction is employed. For the simplest single-bit error detection code, $n_r \geq \log(n_d) + n_d$.* (This assumes that the address is implicit in the position of the data bit within the bitstream. Other viable alternatives, like content-addressable memory, are not considered here.)

There are many possible ways to format the data bits ($d_{i,j}$, $0 \leq j < n_d$) and the error-correction bits ($e_{i,j}$, $0 \leq j < n_r - n_d$) in the bitstream for addresses $2^{n_a} > i \geq 0$. Figure 2 shows two possibilities for the trivial example ($n_r = 5$, $n_d = 2$). The interleaving of error-correction bits within data bits simplifies the design of the hardware. The boldface bits show an example word whose address is $i = 2$. The example on the left places all the bits of the word consecutively within the bitstream, which could cause error-correction to fail if there is more than a momentary interruption in communication. The example on the right places each bit as far away as possible from adjacent bits in its word, allowing for prolonged interference in the bitstream to exhibit only single-bit errors in multiple words, a situation that the simple error-correction scheme used here can deal with. Of course, no error-correction scheme can deal with total communications loss.

$d_{3,1}\ e_{3,2}\ d_{3,0}\ e_{3,1}\ e_{3,0}$ P	$d_{3,1}\ \mathbf{e_{2,2}}\ d_{1,0}\ e_{0,1}$ P
$\mathbf{d_{2,1}\ e_{2,2}\ d_{2,0}\ e_{2,1}\ e_{2,0}}$ P	$e_{3,0}\ \mathbf{d_{2,1}}\ e_{1,2}\ d_{0,0}$ P
$d_{1,1}\ e_{1,2}\ d_{1,0}\ e_{1,1}\ e_{1,0}$ P	$e_{3,1}\ \mathbf{e_{2,0}}\ d_{1,1}\ e_{0,2}$ P
$d_{0,1}\ e_{0,2}\ d_{0,0}\ e_{0,1}\ e_{0,0}$ P	$d_{3,0}\ \mathbf{e_{2,1}}\ e_{1,0}\ d_{0,1}$ P
	$e_{3,2}\ \mathbf{d_{2,0}}\ e_{1,1}\ e_{0,0}$ P

Figure 2: Possible Bitstream Formats

Both alternatives include sector pulses (shown as P), analogous to those on disk drives, to allow the VISC hardware to resynchronize its internal counter if communication is interrupted.

In either alternative of Figure 2, the worst-case word access time is roughly t. Regardless of how the bits are formatted; the constraint that must be satisfied is that the $n_d 2^{n_a}$ bits required by the architecture must fit within the redundant, allocated bitstream:

$$r/a{\cdot}n_d{\cdot}2^{n_a} < n_h{\cdot}f{\cdot}d/c \qquad (1)$$

The parameters d and n_h are the those giving the designer some freedom; the others are more constrained by architecture, technology, power and physics.

* Note that all logarithms in this paper are base-2 logarithms, the 'binary log', $\log_2(\cdot)$.

Relation (2) reveals the expected logarithmic relationship mapping distance, hops and other parameters to a bound on address bits.

$$n_a < \log(d) + \log(n_h) + \log(a) + \log(f) - \log(c) - \log(r) - \log(n_d) \quad (2)$$

The energy to transmit one bit successfully between probes is proportional to the square of the interprobe distance. Since the number of VISC bits transmitted per second is af, the minimum transmission power is proportional to afd^2. We assume that the VISC and repeaters operate as a Globally Asynchronous Locally Synchronous (GALS) system, wherein onboard oscillators operate at a frequency nominally n_f times faster than the transmission bitrate. During a typical period of about n_f cycles, the onboard receiver is able to distinguish between the receipt of a first valid bit (0 or 1), a quiescent period, and the receipt of a second valid bit. This three-state channel (0, 1 and quiescent) allows asynchronous operation. The minimum power requirement is the sum of static power P_S (which includes any payload sensors), CPU (logic) power P_L (which includes the receiver) and transmission power P_T. Assuming we may ignore the static power, the power required per bit, X, is

$$P = P_S + (P_L n_f + P_T ad^2) \cdot f \quad (3)$$

$$X = (X_L n_f/(ad) + X_T d)/n_h \quad (4)$$

For any reasonable interprobe distance, the cost of the serial logic will be insignificant, even though it operates n_f times faster than the transmitter. Assuming we may neglect the digital logic, the on-probe power needed per bit is proportional to d/n_h.

A successful DRIFT deployment of MEMS-enabled probes inherently needs power supplies and transmitters able to communicate some minimum distance, even if the VISC concept were not being used. Given this as d, we may choose n_h to satisfy any selected VISC n_a in (2). In the simplest single-threaded implementation, only one probe has a VISC CPU, and $(n_h - 1)$ probes are repeaters. Only $t^{-1} \approx 10^4$ non-memory-reference instructions per second can execute on average, and memory reference instructions might be slower. This might seem a poor utilization of the digital logic.

A more efficient alternative would take a multi-threaded approach, where each of the n_h probes has a VISC CPU, which would hold the state of one thread. As the bitstream rotates around the n_h probes, each of the threads will have the opportunity to access or update bits in the program and data shared by all the threads, allowing on the order of $n_h/t \approx 10^6$ instructions per second.

6 SPACECRAFT DESIGN

A VISC computer is not an immediately practical proposition, but most of the novel aspects are already under active research in ongoing projects. For reference, we consider a VISC node as a stripped-down evolution of a single spacecraft-on-a-chip concept. The 'SpaceChip' is a proposed monolithic spacecraft based on commercial CMOS, with a target unit cost of $1000; a preliminary design is described by Barnhart *et al.* [2] with consideration to the various system budgets of a conventional satellite.

Adopting this approach, the basic subsystems usually of interest are the structure, electrical power system (EPS), communications, propulsion, attitude control (ACS), thermal control (TCS) and command and data handling (CDH). We consider a potential deployment of VISC to be provided by a carrier spacecraft, which releases each node along a known trajectory with some appropriate Δv. The next assumption is that the payload capability of any given node does not require active pointing: either orientation is irrelevant, or a stable orientation is established through ballistics or the deployment mechanism. Propulsion and ACS are therefore removed from our monolithic conceptual spacecraft.

Power (EPS)

Power distribution, regulation, and control are often implemented in CMOS today, and introduce no significant complexity to the node design. Power generation is a different matter however, with external sources assumed to be available. Monolithic (same die) photovoltaics are impractical for typical

CMOS processes since the efficiency of silicon solar cells is very low compared to more typical materials used for conventional cells; GaAs or GaInP offer much greater power conversion efficiency but more complex fabrication. On-die energy storage (*e.g.* using large capacitors) is also discarded as requiring prohibitively large die areas even for high capacitance options such as bipolar SiGe.

As reported in [2], communication power requirements increasingly dominate power budgets as other contributions are reduced; the free-space loss on a link places lower bounds on energy-per-bit regardless of the sophistication of the remainder of the system. The heaviest non-communications subsystem of the SpaceChip is the payload, an 80μW CMOS imager. In the limit that ACS, propulsion, power management and data handling start dissipating negligible microwatts, communications energy accounts for almost 100% of the power budget.

Link Budget

Recent experiments have shown that atomic-scale carbon nanotubes can act as radio receivers and transmitters [18] and RF MEMS [11] are already available to link the microprobes together. Temporarily ignoring the important loss factors, receiver gain, and any coding gain in the link budget, a 1μW (EIRP = -60dBW without antenna gain) transmission over 1km already creates a receiver challenge in detection of a signal at -161dBW. If we are free to tolerate low bit rates (512bps), implement small (1mm) antennas, and assume a relatively low system noise temperature (400K), an E_b/N_o of approximately 16dB is required.

7 MISSION DESIGN

There are many topological and geometrical alternatives for deployment of DRIFT probes. To estimate the number of probes, n_p, needed to completely cover a planet of radius R at some altitude h using nearest neighbor interprobe communication at distance d, we will assume the area covered by one probe communication range will contain n_h probes. Thus,

$$n_p = 4n_h(R+h)^2/d^2 \quad (5)$$

For example, for a mission to Venus ($R = 6{\cdot}10^3$km) with simple repeater VISCs ($n_h = 2$) having a range of around $d = 100$ km at (very low) orbit altitude $h = 300$ km, would require over $n_p = 31750$ probes to cover the upper atmosphere. Although the most obvious concept would provide universal connection between all probes via relay from one group of n_h probes to a similar neighbor group and so on, an easier alternative would be to isolate groups of n_h probes, and use a conventional relay spacecraft to interrogate each group's memory contents as it travels in a relatively tight (maximum separation $\leq d$) formation.

In order to provide universal connection, it might be helpful to provide two (or more) transmission power levels. The lowest power would allow transmission with bandwidth f over a distance d as described earlier. Higher power levels would transmit less information over longer distances. Part of the unused $(1 - a)$ raw link capacity could be used to communicate between groups of probes in a way that would not interfere with their memory contents. This requires gateway probes that are close enough to eavesdrop on the low-power local VISC memory traffic and high-power communication from distant VISCs. From this, the gateway can choose the proper time to forward the information to the local VISC, allowing for a linear network of arbitrary length, without the need for a mothership.

8 INTERPLANETARY AND INTERSTELLAR EXTENSIONS

The preceding section assumed that a mothership a) transports the MEMS probes from Earth; b) deploys the probes near the target planet, and c) forwards communication from the VISC/DRIFT network back to Earth via a conventional high-power transmitter. A more radical implementation of the VISC concept would eliminate the mothership, and instead establish an unbroken train of probes that forms a relay network from Earth to the planet and back. To enable this vision, we assume an electrostatic or electromagnetic accelerator catapults each VISC component on a path so that it will eventually fly near the planet. Such accelerators are quite

possibly feasible with current technology because the mass of each probe could be on the order of nanograms, and therefore the energy cost of acceleration is trillions of times less than to accelerate a macroscopic kilogram-size probe.

Since gravitational interaction is less significant than electrostatic acceleration at this scale, more-or-less straight-line trajectories should be possible. Assume the distance from near earth to the planet (including compensation for relative motion) is D and the train of probes travels at velocity V relative to Earth. To maintain an interprobe distance of d as assumed earlier, probes must be launched from near Earth orbit at a rate of $F = V/\gamma d$, where the Lorentz factor $\gamma \approx 1.0$ for modest velocities.

From the perspective of Earth, the total mission time, $T = T_1 + T_2 + T_3$, involves three phases: the time for the first probe to arrive, $T_1 = D/V$; the time for gathering, processing and compressing data, $T_2 = n_v t\gamma$, occupying n_v VISC memory cycles; and the time for transferring the data back, $T_3 \approx D/c$. To maintain an uninterrupted train during this time requires n_t probes:

$$\begin{aligned} n_t &= T \cdot F \\ &= D/(d\gamma)\cdot(1 + V/c) + n_v n_h V/c \end{aligned} \tag{6}$$

Before embarking upon an example, it ought to be apparent that n_t is an extremely large number of devices to form into a single mission concept. Although it should be possible to launch such a train of probes to objects in the solar system at velocities similar to that of chemical rockets, the radical potential of the idea becomes evident if one considers accelerating the probes to relativistic velocities (say $V = 0.5c$, $\gamma \approx 1.15$) as would be needed for an interstellar mission (consider the usual example of α-Cen[†], with $D \approx 4.1\cdot10^{13}$ km).

With a launch frequency of about 1500 probes per second, giving $d \approx 100$ km, we would need about $n_t \approx 5\cdot10^{11}$ probes. To put this seemingly huge number in perspective, again assuming $m = 10^{-9}$g, the mass of the entire train for the 12-year mission would be approximately 0.5 kg, requiring a total of $3\cdot10^9$ kJ to launch, which we estimate to cost around €3m/year.

The basic design of the VISC network does not change much, regardless of train velocity, because the relative velocity of adjacent probes will be nearly zero. Several vaguely futuristic proposals, like the 'starseed launcher' [19] and 'astrochicken' [20] have envisaged this kind of lightweight interstellar exploration, but not to the extreme case wherein the train of probes extends all the way back to Earth.

Those earlier proposals are not even close to feasible with current technology because they require almost organic reorganization (as the name 'astrochicken' implies) of the probes upon arrival (scavenging for resources, building an antenna powerful enough to contact Earth *etc.*); currently we do not know how to program a mission *a priori* with such complexity, starting with 'starseeds' as tiny as we are discussing.

The novel innovation proposed here is that a massive swarm of very simple probes eliminates the need to create a self-replicating organism. Instead we have a large but simple computer network that happens to store most of its bits in the actual photons that link it back to Earth.

Unlike all earlier space missions that use a powerful transmitter to emit an expensive cone (with volume proportional to D^3) of RF energy all the way back to Earth, our novel VISC train uses a massive number of weak repeaters to make an economical cylindrical thread (with volume proportional to d^2D) of RF waves that acts like a network cable to Earth. The power requirements for the earlier 'starseed launcher' proposal [19] were supposed to be met from an Earth-based microwave laser directed along the path of the probes. The same approach could be usable here, but with much lower power requirements. Alternatively, probes may be able to capture energy from their interaction with ions and/or magnetic fields.

[†] The alpha Centauri binary system

9 CONCLUSIONS AND NEXT STEPS

We have proposed a way of thinking about interplanetary probes that involves cost saving through mass production of simple things: tiny digital logic, MEMS micro-spacecraft systems and carbon nanotube RF transceivers. Interplanetary probes are feasible only with onboard computers that allow them to operate autonomously, yet the size and speed of computer required is modest enough to consider using the bit-serial approach proposed here.

The presented VISC concept eliminates tangible memory, and actually stores information using RF communication links. Error-correction helps overcome the unreliability of such links and multi-threaded software helps overcome the inherent slowness of bit-serial processing. We have performed preliminary validation of these concepts using simulation and synthesis of Verilog models for a simple VISC. We have proposed a distributed, redundant and inherently fault tolerant (DRIFT) approach to deploying a large number of such nanoprobes into the atmosphere of a planet to provide an observation platform for extended collection, compared to typical entry probes, of simple physical parameters. Complex payloads have not been considered in detail: a generic approach that supposes a variety of advanced MEMS devices could be integrated is assumed.

Solar-powered MEMS technology at the 1000 bit/mm/mg scale [11] already exists that can implement the concepts described here. We anticipate that advances in nanotechnology may allow large numbers of simple devices to someday form planet-wide and interstellar networks.

VISC/DRIFT may well not emerge to be the preferred embodiment for exploration nanoprobes, but it illustrates the possibilities via the paradigm shift of thinking of nanoprobe space exploration as a variant of computer science and how the new emphasis in computer science on cyber-physical systems can be enabling for grand space exploration.

Three physical or technological challenges are identified which are worthy of further research and development:

1. Establishing viable communication links over reasonable distances (at least 10^2 m for simple missions, or 10^6 m for relativistic missions) using nanotube transceivers attached to simple digital state machines
2. Integrating nanotube transceivers with MEMS or conventional digital chips
3. Using electromagnetic interaction of an accelerated probe and its environment to harvest some of its kinetic energy for powering logic and nanotube transceivers.

These developments are left as *exercises to the reader*, while further development of VISC and its other potential applications is undertaken.

10 REFERENCES

[1] D.J. Barnhart, *Very Small Satellite Design for Space Sensor Networks*, 2008.

[2] D. Barnhart, T. Vladimirova, and M. Sweeting, *Very-Small-Satellite Design for Distributed Space Missions*, Journal of Spacecraft and Rockets, vol. 44, Dec. 2007, pp. 1294-1306.

[3] H.R. Shea, *Reliability of MEMS for Space Applications*, in Reliability, Packaging, Testing, and Characterization of MEMS/MOEMS V, Proceedings of SPIE, 2006, pp. 61110A-61110A-10.

[4] M. Eiden, *Experience with Research and Development of MEMS for Space Systems – At the Edge of Revolution*, Proc. NanoTech 2002, Houston TX, AIAA, 2002.

[5] T. George, *Overview of MEMS/NEMS Technology Development for Space Applications at NASA/JPL*, Proceedings of SPIE, vol. 5116, 2003, pp. 136-148.

[6] J. Blamont and R.Z. Sagdeev, *The VEGA Mission*, *Naturwissenschaften*, vol. 71, Jun. 1984, pp. 295-302.

[7] J. Lebreton and D. Matson, *Huygens: Science, Payload and Mission (SP-1177)*, ESA, 1997.

[8] N. Wells, A. Ballard *et al.*, *Atmospheric Microprobes for Venus: A Preliminary Probe Design and Localisation Method*, Proceedings of the 8th ESA Workshop on Advanced Space Technologies for Robotics and Automation (ASTRA 2004), ESTEC, The Netherlands, 2008.

[9] M.V. Wilkes, *Computers Then and Now*, J. ACM, vol. 15, 1968, pp. 1-7.

[10] CCSDS 100.0-G-1: *Telemetry - Summary of Concept and Rationale*, Green Book Issue 1, Dec. 1987.

[11] R. Johnson, *Spacecraft-on-a-Chip to Harness MEMS*, MEMS Microfluidics and Microsystems Executive Review, April 2010.

[12] W.D. Hillis, *The Connection Machine*, MIT Press Series in Artificial Intelligence, Apr. 1986.

[13] M. Quinn, *Parallel Programming in C with MPI and OpenMP*, McGraw-Hill Inc., 2004.

[14] *CUDA Reference Manual*, Version 2.3, Jul. 2009.

[15] M. Arnold, *Verilog Digital Computer Design*, Prentice Hall, 1999.

[16] M. Arnold, P. Vouzis, and J. Cho, *Bitstream Efficiency of Field Programmable One Hot Arrays*, Proc. ISVLSI, Greece, July 2010.

[17] C. Brothers and D. Alexander, *Radiation Hardening Techniques for Commercially Produced Microelectronics for Space Guidance and Control Applications*, Guidance and Control 1997, Advances in the Astronautical Sciences, vol. 94, 1997, pp. 169-180.

[18] H. Peng, C. Chang *et al.*, *Microwave Electromechnical Resonador Consisting of Clamped Carbon Nanotubes in Abacus Arrangements*, Physical Review B, vol. 76, 2007.

[19] Institute of Atomic-Scale Engineering, "The Starseed Launcher," 1996.

[20] G.L. Matloff, *Deep-Space Probes*, Springer, 2006.