

Long Range Weather Prediction III: Miniaturized Distributed Sensors for Global Atmospheric Measurements

E. Teller, C. Leith, G. Canavan, L. Wood

This article was submitted to
26th International Symposium on Planetary Emergencies, Erice,
Italy, August 20-24, 2001

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

November 13, 2001

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

This report has been reproduced directly from the best available copy.

Available electronically at <http://www.doe.gov/bridge>

Available for a processing fee to U.S. Department of Energy
and its contractors in paper from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

Available for the sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/ordering.htm>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

Long Range Weather Prediction III: Miniaturized Distributed Sensors For Global Atmospheric Measurements*

Edward Teller^{1,2}, Cecil Leith², Gregory Canavan³ and Lowell Wood^{1,2}
Hoover Institution, Stanford University, Stanford CA, U.S.A.¹, and
University of California Los Alamos³ and Lawrence Livermore² National Laboratories,
Los Alamos NM and Livermore CA, U.S.A.

ABSTRACT

We continue consideration of ways-and-means for creating, in an evolutionary, ever-more-powerful manner, a continually-updated data-base of salient atmospheric properties sufficient for finite differenced integration-based, high-fidelity weather prediction over intervals of 2-3 weeks, leveraging the 10^{14} FLOPS digital computing systems now coming into existence.

A constellation comprised of 10^6 - 10^9 small atmospheric sampling systems – high-tech superpressure balloons carrying early 21st century semiconductor devices, drifting with the local winds over the meteorological spectrum of pressure-altitudes – that assays all portions of the troposphere and lower stratosphere remains the central feature of the proposed system. We suggest that these devices should be active-signaling, rather than passive-transponding, as we had previously proposed only for the ground- and aquatic-situated sensors of this system.

Instead of periodic interrogation of the intra-atmospheric transponder population by a constellation of sophisticated small satellites in low Earth orbit, we now propose to retrieve information from the instrumented balloon constellation by existing satellite telephony systems, acting as cellular tower-nodes in a global cellular telephony system whose 'user-set' is the atmospheric-sampling and surface-level monitoring constellations. We thereby leverage the huge investment in cellular (satellite) telephony and GPS technologies, with large technical and economic gains.

This proposal minimizes sponsor forward commitment along its entire programmatic trajectory, and moreover may return data of weather-predictive value soon after field activities commence. We emphasize its high near-term value for making better mesoscale, relatively short-term weather predictions with computing-intensive means, and its great long-term utility in enhancing the meteorological basis for global change predictive studies.

We again note that adverse impacts of weather involve continuing costs of the order of 1% of GDP, a large fraction of which could be retrieved if high-fidelity predictions of two weeks' forward applicability were available. These $\sim \$10^2$ B annual savings dwarf the $< \$1$ B costs of operating a rational, long-range weather prediction system of the type proposed.

* Prepared for presentation at the 26th Symposium on Planetary Emergencies, Majorana Centre for Scientific Culture, Erice, Italy, 20-23 August 2001. Work performed in part under the auspices of the U.S. Department of Energy, in the course of Contracts W-7405-eng-36 and -48 with the University of California. *Opinions expressed are those of the authors only.*

Introduction and Summary. Weather significantly impacts a very wide variety of human undertakings, including all outdoors recreational, economic and defense activities. We have previously estimated its adverse impacts to have aggregate costs of the order of one percent of GNP.¹ Thus, improved ability to forecast weather – over longer distances in time, in higher geographic resolution and with greater fidelity-of-forecasting – and thereby to mitigate its disruptions may be expected to lead to aggregate economic benefits of several tens of billions of dollars per year in the U.S. alone, and several times this world-wide.

Computational Capability Gains Emphasize Geophysical Data Shortages. Improved global measurements of the state of the atmosphere, correspondingly enhanced physical models, and advanced computational capabilities all would be required to support such rational, enhanced predictive capabilities. It appears likely that the requisite computational capabilities – arithmetic engines with $\geq 10^{14}$ Floating-Point Operations per Second, or 0.1 petaFLOPS capabilities – will come into existence during the next few years for a wide variety of other digitally-based physical simulation purposes, ranging from nuclear explosions (e.g., USDoE ASCI, which has just taken delivery on a 1.2×10^{13} FLOPS system) to folding proteins (e.g., IBM ‘Blue Gene’, which is projected to attain 10^{15} FLOPS capability).

The progress rate-limiting consideration thus appears to be improved knowledge of the time-dependent status of the atmosphere (and, to lesser extents, that of the upper oceans and the land-surface) and the enhanced physical models that may be expected to naturally arise swiftly from such improved knowledge. In this paper, therefore, we continue to focus on swift attainment of large improvements in this geophysical knowledge base.

Improving Global Change Forecasting. We note that ‘new’ atmospheric science learned in the course of greatly improving measurements of atmospheric conditions could also remove impediments that have reduced the utility of weather models for climate studies and of climate analyses for weather forecasting. Such amelioration of this fundamental logical disconnect would improve the utility of both types of modeling, as well as enhance the overall level of confidence in the corresponding predictions regarding long-term effects of anthropogenic forcing of the atmosphere, land and oceans.

The research program needed to evaluate the usefulness of a large population of rather uniformly distributed sensors for global fluid-geophysical measurements could be based on a modest extension of current technology. It would also provide a more quantitative assessment of the eventual impact of such measurements on improved weather and climate forecasting.

Central Concept. We continue to suggest that global atmospheric measurements are most aptly performed by small balloons designed to float at various altitudes of meteorological significance and which carry high-technology means for conducting *in situ* measurements of local position and wind velocity, as well as ambient temperature, water-vapor density and other such other physical variables as may be deemed essential for the specification of the Earth’s atmosphere and calculation of its time-evolution. These balloons act as Lagrangian trace particles, with their motions providing a direct, accurate measurement of local wind velocity, which is arguably the most important meteorological variable for accurate weather predictions, but which is difficult to measure or infer accurately remotely. Carried-along sensors attached to

¹ E. Teller et al., “Long-Range Weather Prediction Enabled by Probing of the Atmosphere at High Space-Time Resolution,” Univ. of Calif. Lawrence Livermore Nat’l. Lab. Report UCRL-JC-131601 (August 5, 1998), presented at the 23rd Symposium in 1998.

the balloons can measure local atmospheric pressure and temperature, insolation, water-vapor and -droplet densities, and trace constituents.

A modest constellation of such balloons, more or less uniformly distributed and numbering about 3×10^4 (and perhaps having a cost-to-create-at-the-margin of a few million dollars) could provide measurements at the horizontal resolution $D = 200\text{-}300$ km – characteristic of most contemporary general circulation models. In general, the number of balloons required increases as $1/\{D^2 \ln D\}$, where D is the mean horizontal spacing between the balloon micro-airships; the logarithmic term accounts approximately for the greater vertical resolution required as the horizontal sampling-length shrinks. As we will discuss in more detail below, the resulting weather-prediction times only increase logarithmically with the horizontal resolution D . However, a comparatively modest number of balloons ($\sim 10^5$) could double the time-intervals over which predictions are accurate, relative to the best present-day ones, and data returned from an eminently practical number of balloons ($10^8\text{-}10^9$) could support atmospheric measurements at the theoretical limit of the weather's predictability.

We note that detailed results of our collaborators, which we presented at this meeting two years ago,² offered a first-order demonstration-in-simulation that such Lagrangian tracers actually remain quite uniformly distributed in GCM model-studies over multi-month intervals. We also invite recollection of the Global Atmospheric Sampling Program (GASP) results of several decades ago, in which the mean lifetime of superpressure weather balloons in the troposphere was demonstrated to be of the order of six months, with icing-out in tropical storms being the dominant loss mechanism.

We find it remarkable that a constellation comprised of 10^9 such atmospheric monitoring devices – which would apparently have a marginal creation cost of the order of \$1/unit – could sample the atmosphere at the currently-perceived spatial resolution limits of its intrinsic, chaos-limited predictability, which apparently is between 2 and 3 weeks in future-time, and that an exponentially more modest-sized constellation of 10^7 such devices likely could provide atmospheric sampling data sufficient for predictions of 2 weeks' validity in future-time. We now discuss such scaling considerations in a little detail.

Constellation sizes can be estimated geometrically, as the balloons are intended to give roughly uniform coverage over the whole Earth. If the average distance between balloons is D , each balloon 'covers' an area $\pi(D/2)^2$. The area of the Earth's surface is $4\pi R_e^2$, R_e the Earth's radius. Thus the number of balloons required to give n vertical measurements with a mean horizontal separation of D is

$$N^* \approx nN = n\{4\pi R_e^2/\pi(D/2)^2\} = n\{(4R_e/D)^2\} \quad (1)$$

For a single layer of sampling, e.g., at the 500 mb pressure-altitude, uniform horizontal resolution with $D = 1,000$ km would require a constellation population N of 650 balloons; while a D of 300 km would involve 7,200, and one of 10 km would take 6.5 million. Adding vertical measurements that increase logarithmically with the horizontal spatial-sampling frequency – e.g., from 2 vertical-sampling layers at 3000 km horizontal resolution to 10 such layers at 3 km horizontal resolution – would increase the constellation-size for 1,000, 100 and 10 km mean horizontal spacing to about 2,100, 390,000, and 56 million, respectively. These data are among

² E. Teller et al., "Long-Range Weather Prediction And Prevention Of Climate Catastrophes: A Status Report," Univ. of Calif. Lawrence Livermore Nat'l. Lab. Report UCRL-JC-135414 (August 18, 1999), presented at the 24th Planetary Emergencies Symposium in 1999.

those presented in Figure 1. As the balloons are estimated to have average lifetimes on the order of a year, maintaining any of these constellations would require launching roughly these numbers of balloons annually.

Prediction time-intervals vary with the resolution of meteorological measurements, particularly the air's vector velocity, at the initial time for the prediction. The Earth's large-scale flows are essentially two-dimensional turbulence; hence, they are only partially predictable. There are no fundamental theories for their evolution; current numerical predictions are simply brute-force integration of approximate equations of motion for the largest scales with some prescription for the effects of sub-grid scales. Currently, imperfect knowledge of initial conditions on smaller distance-scales degrades the accuracy of such integrations at about 4-5 days into future-time. However, turbulence models have advanced to the point where they can make semi-quantitative estimates of the time required for uncertainty about the time it takes for uncertainty about the smallest scales to propagate to larger scales and contaminate their prediction. One such model, the Lagrangian History Direct Interaction Approximation, can estimate the rate of error propagation and is thought to be accurate to within numerical parameters on the order of unity.

Such models confirm the intuitive notion that the time-scale for contamination of the information at a given scale is roughly the circulation time for turbulent eddies of that size, i.e., the time for an eddy of size D to rotate about its axis. The result is that each factor-of-two improvement in resolution should add about one day to the future-time over which the large scales that drive mesoscale weather are predictable. Given resolution of a few kilometers, the ultimate predictability of the atmosphere is estimated to be from 14 to 21 days, depending on the strength of the two-dimensional turbulence in the scales affected. A reasonable representation of the predictability estimates, in days of future-time T , for moderate levels of turbulence is

$$T \approx 3.3 + 1.9 \ln(R_e/D) \text{ days,} \quad (2)$$

which gives predictability times of ~4.7 days at $D = 3000$ km, 9.1 days at 300 km, 13.5 days at 30 km, and 18 days at 3 km.³ Equation (1) can be solved for D and substituted into Equation (2) to produce

$$T \approx 3.3 + 1.9 \ln(\sqrt{N/n}/16) \text{ days} \quad (3)$$

Figure 1 plots total constellation size N^* (i.e., nN), number of vertical layers n , number per horizontal layer N , mean horizontal spacing D , and normalized costs for constellation creation/maintenance for learning-curve power-coefficients of 0.9, 0.8 and 0.7, which is believed to span the range of likely interest, all as functions of future-time of prediction validity T . It indicates that the current prediction time of 4-5 days could be supported by about 200 balloons, 9 days by 30,000 balloons, and 14 days by 10^7 balloons. These numbers quickly become large with increasing future-time of prediction validity, although they're certainly not excessive in overall cost-benefit terms. Fundamentally, it appears that the number of balloons in the atmospheric-sampling constellation might aptly be increased progressively, in synchrony with capabilities for data return and for computational analysis, with the total size of the

³ G. Canavan, "Value of Global Weather Sensors," Los Alamos National Laboratory Report LA-UR-99-0018 (1999).

constellation deployed being determined by incremental costs versus incremental societal benefits realized for longer-duration, higher geographical accuracy weather predictions.

The atmospheric-sampling interval required is roughly the Lagrangian turnover time for eddies of the size of the separation between balloons. Smaller sizes are noise, which must be averaged out. Larger sizes produce the mean motion that only advects the eddies. In fixed-position sensors, this produces rapid time variation. Advection of a frozen field of eddies of spatial frequency $k = 2\pi/D$ by a large scale random velocity U produces a time-varying signal e^{ikUt} . Its correlation function is $e^{-(kUt)^2}$, so measurements decorrelate on an Eulerian time T_E

$$T_E \approx 1/kU \quad (4)$$

For $D = 100$ km, $k \approx 6 \times 10^{-5}/\text{m}$, so for $U = 30$ m/s, $T_E \approx 1/(6 \times 10^{-5}/\text{m} \times 30 \text{ m/s}) \approx 500$ s, so that sensors distributed with this mean horizontal spacing would have to make measurements in a time short compared to this.

In a three-dimensional Kolmogorovian inertial cascade, the turnover time for an eddy of size D is $t \sim D/u$, where $u \sim (\epsilon D)^{1/3}$ is the eddy's Lagrangian velocity, and $\epsilon \sim U^3/L$ is the rate of energy cascade from the largest eddies of velocity U and scale L through smaller ones. However, in the large-scale, two-dimensional portion of the spectrum which characterizes the actual atmosphere, the semi-conserved quantity is enstrophy, or specific angular momentum, which cascades at a rate $\sigma \sim (U/L)^2/(L/U) \sim (U/L)^3$ to larger scales. σ is not known with precision, but it has the dimensions of $1/\text{time}^3$, so inertial range eddy velocities must scale as $u \sim (\sigma D^3)^{1/3} \sim D\sigma^{1/3}$. Eddies of size D are 'contaminated' on their turn-over time-scale, which is their Lagrangian time

$$T_L \sim D/u \sim D/\sigma^{1/3}D = 1/\sigma^{1/3}, \quad (5)$$

which is independent of eddy size. Hence, the total predictability time is the number of scales times the time-scale to corrupt each, which is $1/\sigma^{1/3}$. For resolution $D = 3$ km, the ratio of scales is $\approx 6,371/3 \approx 2000$, so the number of scale-doublings is $\approx \log_2 2000 = 11$. Thus, the maximum predictability time of $T \approx 18$ days results from 11 successive contaminations, each of which $1/\sigma^{1/3}$, so $T \approx 11/\sigma^{1/3}$, and $\sigma^{1/3} \approx 11/T \approx 0.55/\text{day}$. The local velocity for $D = 1,000$ km is $\approx D\sigma^{1/3} \approx 10^6 \text{ m} \times 5 \times 10^{-6}/\text{s} \approx 5$ m/s, and that for 100 km ≈ 0.5 m/s, but both must be sampled on time-scales short compared to $1/\sigma^{1/3} \approx 2$ days in order to properly capture local dynamics. This time is longer than the Eulerian measurement time of Eq. (4) by a factor

$$t_L/t_E \sim (D/u)/(D/U) \sim U/u \sim \sigma^{1/3}L/\sigma^{1/3}D \approx L/D, \quad (6)$$

which for 100 km eddies is a factor of $\approx 6371/100 \approx 63$.

We therefore take the required atmospheric sampling time-interval to be $\sim 10^{-2} \times 1/\sigma^{1/3}$, or $\sim 10^3$ seconds for the smallest distance-scales. A data-frame of ~ 100 bytes, supplemented by headers, error-detection/-correction syndromes, etc., would involve data-packets of $\sim 10^3$ bits in total length, so that each atmospheric sampling platform would return a time-averaged data-rate of 1 bit/second. A constellation of $\sim 10^6$ platforms thus would have a total data return-rate of 10^6 bits/second. Over the IRIDIUM satellite telephony constellation, for instance, this would amount to $\sim 10^4$ bits/second-satellite, which is of the order of 1% of available channel capacity.

It's therefore clear that atmospheric-sampling constellations of the scales presently contemplated would utilize only tiny fractions of the capabilities of even extant satellite telephony systems. The constellations of ocean- and land-surface-sampling platforms would have comparable data-return rates as their atmospheric cousins but would be substantially less populous, thereby adding only incrementally to the overall loading of the telecommunications sub-system just estimated.

Atmospheric Measurement Technologies. In our initial study of this subject, presented at this Symposium three years ago, we proposed the use of entirely passive atmospheric transponders, maximally simple devices that merely changed the frequency- or polarization-dependence of their reflectivities as the local temperature and relative humidity changed. The time-dependent positions of these Lagrangian particles were sensed by probe laser-beams emitted by a constellation of small satellites in low Earth orbit, with pairs of position-sensings made at the ends of a small, known time interval providing the vector velocities of movement. The local temperature and humidity of course were 'read out' during each such probing event, both to moderate (but necessarily not really high) accuracy.

While we continue to believe that such a system is eminently feasible from a technical standpoint, its relatively large initial cost – associated with creation of the LEO constellation of transponder-interrogating smallsats – is now recognized as a significant programmatic impediment. We therefore have considered alternative approaches to the same basic atmospheric sensing capability with substantially lower 'up front' costs. Our thoughts have naturally turned toward leveraging commercial off-the-shelf capabilities, ones that can be purchased, rather than developed *ab initio*. In this spirit, we're naturally attracted by 'free' capabilities, ones that may be borrowed or shared with others, e.g., Global Positioning System (GPS)-based sensing of local position and velocity.

Sensing Of Atmospheric Properties. We presently believe that the cost-efficiency optimum of any system architecture directed to present ends centers on somewhat larger-than-previous superpressure balloons deployed in the 'meteorological' portions of the atmosphere – ones of 30-100 cm diameter having individual volumes of a few dozen to several hundred liters deployed over the first ~15 km in altitude above sea-level. Each of these micro-airships would carry as a minimal, 'asymptotic' payload a modern top-end cellular telephony chip complete with embedded GPS receiving and control-processor circuitry which supervises various micro-miniaturized environmental sensors, along with a small battery and a photovoltaic power-supply, with a total payload-system mass of the order of 10 grams. The notional layout of a monolithic-silicon implementation of this system is depicted in Figure 2. When at their various operational altitudes – i.e., when fully-inflated – such superpressure balloons could very readily and mass-efficiently deploy a sub-micron thickness of GaAs solar photovoltaic array and the set of L-band antennae needed for multi-channel GPS reception, as well as the cellular telephony antenna, all on the same thin plastic balloon-wall ($\sim 10^{-2}$ cm). All of these electromagnetic structures could be implemented as either thin semiconductor sheets or traces of metal having very modest mass budgets, and yet could source of the order of a single time-averaged watt of electrical power into the secondary battery and balloon-payload, as well as serve as quarter-wave dipole antennae for cellular telephony transceiving and GPS signal-reception.

Maximally-Simplified Data-Return. While in the atmosphere over continental territory which has been sufficiently highly developed to have cellular telephony towers, such devices

could simply "call home" periodically, in order to download their data to the central data-processing facility through the global public switched telephony network (PSTN), thereby providing wind velocity, temperature, humidity, pressure, insolation, etc. data-streams along the space-time trajectory defined by the motion of the atmospheric packet in which each of them is embedded as a Lagrangian particle. Indeed, this "maximally commercialized" mode of data-return would be eminently suitable for initial proof-of-principle studies, as it would return highly useful data-streams at absolutely minimal investment. Quite importantly, its bi-directional character would permit the central data-processing facility to re-program the balloons' embedded digital controllers in a serial number- or location-dependent manner, by uplinking new software-loads or, even more readily, by simply changing key parameters being referenced by the on-board control processor. Each balloon could thus serve as a truly interactive mobile weather-station, dynamically modifying its data-collection and -return behavior in order to optimize the global performance metrics of the main Weather-Prediction Engine, as that Engine's environmental data requirements might vary in time.

However, it appears that this data-return scheme would be useful only over economically highly-developed portions of the Earth's land-mass, and couldn't possibly function satisfactorily when the balloons were carried by air-motion over mountains, deserts, jungles, lakes and oceans, where cellular telephony towers are yet to be found – which is likely to be at least 95% of the Earth's total surface, recalling that only 15% of the Earth's landmass is truly arable, and most of that is yet to be 'blessed' with cellular telephony. Data-streams returned by balloons just coming into range of a cellular telephony tower, e.g., on a western continental margin, thus would be cluttered with "stale data" typically collected dozens to hundreds of hours previously – and thus would have relatively little content of then-current value for weather-prediction purposes. Of comparable importance, ground tower-based cellular telephony is intrinsically incapable of returning geophysical data from most all members of the ocean- and land-based constellations-of-environmental monitors, data that are essential as boundary-condition inputs to the numerical integration procedures which creates the long-range weather forecasts.

Satellite Telephony Data-Return. However, this scarcity-of-ground-based-towers picture is quite likely much too bleak. There are already multiple satellite telephony services in commercial operation, and any one of these having truly global 'reach' could potentially provide world-wide real-time connectivity between any one of our atmospheric-monitoring micro-airships and the main Weather-Prediction Engine. One of these satellite telephony systems, the IRIDIUM one discussed at a Symposium session several years ago, following an interval of financial and organizational turbulence, currently has the U.S. Defense Department as its 'anchor tenant' and presently provides unlimited service to 20,000 DoD-owned user hand-sets around the world. With only a small fraction of its data-handling capacity, IRIDIUM could service the data-traffic of literally a million or of balloon-borne atmospheric monitors, each of which might be expected to return a data-packet of 1 kilobytes length – e.g., ~32 data-frames of ~32 bytes each, taken regularly at tenth-hour intervals – in a few-second-duration, burst-type 'call home' every few hours, comprising a steady-state loading on the IRIDIUM system of 'only' 300 calls, or a half-dozen calls per satellite.⁴ Interestingly, data return could

⁴ Some concern might arise as to whether the substantial mass and volume of contemporary satellite telephony handsets – the smallest of which have ~300 gm masses – are compatible with carriage by micro-airships with total payloads of the order of 10 grams. Our response to this is that these handsets are built to support continuous communication over intervals of more than one hour, and have power supply and heat-dissipation

be scheduled to take place normally only outside of urban areas, so that de-confliction relative to the geographical peaks of commercial traffic-loadings could be attained quite readily; the incremental loading of the system by the constellation of micro-airships thus would be virtually invisible to IRIDIUM system operators (or to other system customers).

Since the U.S. Defense Department is a not-unlikely sponsor or co-sponsor of any American component of any international long-range weather-prediction program, use of the IRIDIUM system for both initial and steady-state data-return purposes is an especially interesting prospect from both technological and economic perspectives, one which could greatly facilitate commencement of operation of the capability which we propose. Of course, there are other commercial satellite telephony systems that offer alternate paths to this initial operational capability, so that this particular programmatic option should be regarded as a quite feasible, low-risk one.⁵

The overall layout and data connectivity of the global environmental monitoring system that we contemplate as a baseline one is indicated schematically in Figure 3. We emphasize that all of the in-space features of this system already exist – and have been in routine operation for years. Only the balloon-borne and surface-emplaced constellations-of-sensors need to be created – and the two fluid-envelope constellations even self-deploy!

Lifetimes And Costs. As already noted, well-designed superpressure balloons have been GASP-demonstrated to have lifetimes in the troposphere of the order of six months – long

means that are very conservatively rated for such service, as customer tolerance of any malfunctions in such premium means of communication is minimal. In acute contrast, we contemplate few-second-duration, burst-type data-communication every few hours, for which the heat capacity of the transmitter chip is adequate heat-sink and the pulse capacity of the tiny albeit high-energy-density (e.g., Li ion) on-board battery is sufficient power-source. In addition, the balloon's size supports deployment of physically optimal antennae, in marked contrast to those of hand-held satellite telephony user-sets. We therefore conclude that gram-scale, silicon-intensive communications systems will be quite capable of sourcing the highly-intermittent, short-duration meteorological data-traffic from the atmospherically-mobile balloon-borne weather-stations that we propose.

⁵ Despite the attractiveness of generic satellite telephony, we consider deployment of a second constellation of balloons in the upper stratosphere as an alternate architecture for the data-return portion of the atmospheric monitoring system of present interest. The purpose of these much larger, much longer-lived superpressure balloons is to serve as the cellular towers of this global telephony system or, from the above perspective, to stand-in for the wireless telephony satellites in orbit.

Employing mesospheric balloon technology which is currently semi-commercially sourced, we contemplate mesospheric data-relay systems with a payload of the order of a few kilograms – about 10^{-3} of the largest payloads currently lofted in such systems – which receive data-calls from the lower-lying atmospheric sampling balloons, out to their local Earth-horizon at a few hundred km distance. The greater size of these mesospheric balloon systems would permit the generation of kilowatt-scales of electrical power, the storage of megajoule quantities of energy, and the deployment of the larger antennae and high-Q filters appropriate for receiving calls from multi-hundred kilometer-distant transmitters having effective radiated powers of a single watt. Each of these systems would then uplink suitably-conditioned data-packets into the global comsat system for return to the main Weather-Prediction Engine.

Of the order of one thousand of these mesospheric “satellites in the atmosphere” would provide world-wide connectivity between the much smaller, much lower-power atmospheric monitoring balloons beneath and the central Weather-Prediction Engine. Each would service between 10^3 and 10^6 monitoring-balloons, at a call-density of 0.1-100 at any given moment, i.e., $10^{-2.5 \pm 1.5}$ of the typical service-level of a single terrestrially-deployed cellular tower.

before modern materials became available. Their lifetimes appears to be limited by catastrophic icing in tropical storms; they circulate about the planet of the order of a dozen times until they are caught in such a storm, take on a relatively heavy load of ice, and fall into the underlying ocean before the ice-burden melts. Our baseline balloon system architecture therefore contemplates a mean lifetime of 0.5 years. [To be sure, we plan to exploit the superpressure feature of our balloons by providing them with an 'emergency lift' capability which may be invoked under control of their embedded processors to expand their displaced volume abruptly by as much as 100% with no increase in mass, thereby giving them the ability to rapidly ascend out of icing regions of storm-clouds and thus escape death-by-icing. Use-as-needed of this feature may extend the average life of our micro-airships by several-fold.]

We thus contemplate a requirement for creating and launching at least 10^6 and no more than 10^9 third- to half-meter-diameter balloons each year – roughly 0.03-30 per second, continually. As we noted three years ago, there is nothing particularly challenging about this requirement, technically, operationally or economically.

Cost Determinants. The experience of the modern chip business suggest that the 'cost-driver' of such balloon-borne weather-station creation likely will be the attachment of the silicon chip+mini-battery module onto the metal bonding-pads on the balloon-wall and the associated functionality-testing; balloon-creation and -inflation to its appointed pressure-altitude seemingly will be next most costly, with procurement of the chip+battery module being third most expensive.

Balloon-associated steps likely will be the more expensive ones, simply because they won't be already-mass-production ones when the program gets seriously underway. On the other hand, disposable cellular phones are just now coming onto the market, with imputed costs of the order of \$5 for the phone that's discarded after a total use-interval of the order of one hour; the associated silicon cost can't be much in excess of \$1 – and this is for unit quantities of an item sold at retail to a rather price-insensitive customer-community. Underscoring this swiftly-emerging trend toward commodity telephony silicon, Motorola, the largest U.S. cellular telephony manufacturer, recently announced that it's commencing to sell its entire cellular telephone component line – including its vaunted "single-chip solutions" – as turn-key packages and kits to any-and-all comers, signaling to industry observers that it's departing a line-of-business which it pioneered in favor of selling mass-market commodity parts, like it's long sold linear semiconductors. Already, cellular phones are going on-sale with built-in GPS receivers and single-button-push means of transmitting the contents of its coordinate registers. We're therefore reasonably confident that total balloon-creation costs-at-the-margin will be \$2-5/unit, with unit marginal silicon costs being of the order of \$1, when micro-airships are created by the millions per year.

System Operating Costs. Our estimate for system operations costs thus is less $\$10^{7.5 \pm 1.5}$ /year and are dominated by balloon-replacement expenses, assuming that cellular telephony costs are kept reasonably in-line. If the system's cellular telephony costs are made to be negligible by such gambits as IRIDIUM usage, if balloons can evade most fatal icing incidents by means such as active pressure-altitude control and thus attain year-scale mean operational lifetimes, and if mass-production cost-scaling can attain a learning-curve power-law index of 0.8, these operating costs for even a maximum-scale system might of the order of $\$10^8$ /year, as is indicated in Figure 1. This represents an order-of-magnitude improvement over our cost-estimates of 3 years ago – and even those far-higher costs were still two orders-of-magnitude less than the gain in economic benefits from long-term, high-fidelity weather

predictions. Much more importantly, however, the buy-in costs to start up the system that we presently propose are the comparatively negligible ones of prototyping powered-silicon-bearing balloons, as compared to those of creating a constellation of smallsats in LEO together with transponder-balloon prototyping.

What Price Distant-Future Knowledge? Finally, we note that the last few days of durability in future-time of weather-prediction accounts for essentially all of the cost of the type of atmospheric monitoring, data-return and computation system that we contemplate. We have estimated above that extension of forecasts of a given level of fidelity by 1 day requires nearly a doubling of resolution in the spatial frequency of atmospheric sampling, i.e., a 3-fold increase in the constellation-size of micro-airship weather stations, with a nearly linear increase in cost. Only the enormous benefit-to-cost characteristics of this forecasting system could possibly justify driving the atmospheric-monitoring constellation size up to the 10^9 -balloon level and on out to the ~18-days in future-time theoretical limit of forecast validity estimated above; indeed, much more detailed analyses of cost-benefit-at-the-margin likely will be required to justify 'buying' the last 2-3 days of future-time forecast validity with 90% of the maximum-scale system's cost.

It is a curious, but potentially quite important, point that the large majority of the maximum-scale system's cost buys knowledge of future weather that seems destined to have the lowest unit value. [Most – but assuredly not all, farmers being an obvious exception – of us care much more about the weather in the next 3-7 days than we do about it on days 14, 16 and 18 from now.]

Some System-Level Considerations. We now address a few system-level considerations, including system operation in the real world, safety, esthetics and benefits.

Operations In The Real World. The balloon lifetime-limiting feature in the lowest 3-5 km of the atmosphere isn't icing, but rather the pilot's bane of CFIT: controlled flight into terrain. Higher-altitude terrain in general and mountains in particular pose as much danger to lower-altitude balloons as they do to powered aircraft, and avoiding crash-landings is the primal motivation for the already-mentioned engineered-in capability employed to evade death-by-icing in major storm systems. After all, every circum-Earth cycle, or roughly every two weeks, all balloons with pressure-altitudes of ≤ 3 -5 km – a third to a half of the constellation – will have to ascend significantly in order to clear the Andes in the Southern Hemisphere or the Himalayas and/or the Rockies in the Northern one. While flying over land, the lowest layers of them will also have to slowly but continuously adjust their pressure-altitudes so as to track reasonably close to the local underlying solid-Earth surface without flying into it.

We contemplate providing such altitude-control capability via tiny, e.g., piezoelectric or rotary-drum, traction-motors attached to filamentary cables within each balloon whose other ends are anchored in the very thin plastic walls. When the balloon's embedded digital controller wishes to reduce the balloon's (pressure-)altitude, the micro-cables are synchronously reeled in by micro-motor action slowly doing the requisite PdV work – about 1 kilojoule's worth, for a 2-fold volume-change in a 20 liter full-volume balloon operating with a base 1 bar pressure-altitude (i.e., a worst-case) – on the balloon's pressurizing-gas and thereby decreasing the balloon's total volume. Conversely, these tensioned filaments may be released (swiftly) under program control when (rapid) ascent is called for. [Other, quite different means of effecting this variable-lift function seem obvious; resistively heating a tiny, insulated-and-

valved reservoir containing a thermally unstable material such as LiH in order to reversibly input ~1 gram of hydrogen gas into the balloon's main buoyancy-generating volume is one such alternative.] Obviously, all of these altitude-control means will be GPS-informed, so that their operation will be entirely automatically implemented by the balloon-resident digital microcontroller, which will carry an all-planet, km-resolution-scale elevation map in a gigabit ROM.

Safety. These balloons that we contemplate are much smaller in size and bear payloads of far smaller mass than do ordinary weather balloons and associated radiosondes. This is the *prima facie* argument for their safety relative to manned aircraft. However, their far greater numbers – which is the source of their socially-valuable weather-predictive power – might give grounds for concern with respect to rare collisions between one of them and a manned aircraft. We note that their contemplated mean density is 0.1-0.001/km³, so that an airplane with a typical cross-section of perhaps 5 m² and a speed of the order of 200 meters/second will have a mean time-to-collision of 10^{10±1} seconds, or 3000-30 years (for constellations of 10⁷-10⁹ balloons, respectively). Even when such a collision occurs, it will have far less adverse consequences than that with a small bird – for few birds that are 'seen' by planes have masses as small as 10 grams. We therefore conclude that there are no significant human safety issues associated with deployment of even large quantities of these gossamer-like micro-airships.

Esthetics. We point out that the micro-airships which we contemplate will be quite invisible from the Earth's surface. The lowest-altitude ones, floating at about 1 km above the local surface, will have an angular subtense smaller than the angular resolving capability of all but the best human eyes. We expect that their underside skins will be optimally mottled in order to further obscure them from Earth-surface view, so as to entirely obviate visual-esthetic concerns.

We are conscious of the need to dispose of fallen balloons in an environmentally-sound manner. The battery-chip module has a mass of the order of 1 gram and the physical scale of a small pebble; covering one-trillionth of the Earth's surface with such readily weather-degraded synthetic pebbles each year seems a very minor environmental impact indeed. When a balloon grounds out – or hits the ocean surface -- we contemplate employing the battery's entire energy-store to cut its ultra-thin plastic skin into cm-scale flakes by electrically-ignited means and to similarly open the protective shell of its battery-chip module to weathering, if the micro-airship descends on land; the latter will swiftly descend into very long-term abyssal storage, if the final descent is into the ocean. While the assessed environmental impacts are not completely negligible, we suggest that they are offset by many orders-of-magnitude by the benefits of long-range, high-fidelity weather prediction, of which enhanced management of large-scale floods is merely a single example.

Benefits. The likely benefits from the proposed globally distributed atmospheric measurement capability fall into three broad categories: science, weather, and climate.

Science Benefits. Atmospheric science is now dominated by uncertainties ranging from predictability to cloud physics, i.e., what will the weather be in a few days (and will we ever be able to improve such predictions) to why does it rain? Remote measurements from the ground, aircraft, and satellites have been too difficult and too expensive to really resolve such basic issues. Detailed knowledge of winds, mixing on all scales, and *in situ* measurements, e.g., of cloud droplet size-density distribution, microparticle density and ion concentration, could provide fundamental insights leading to resolution of such uncertainties.

Weather-Prediction Benefits. Such insight could make it possible to improve weather forecasts from the 4-5 days of the best current techniques, employed under the most favorable prevailing conditions, to one of 2-3 weeks during potentially supportable by more accurate measurements made with much higher space-time resolutions. In the process, the R&D program needed to develop that capability could provide direct information on the atmospheric quasi-turbulent energy spectrum and mesoscale phenomenology needed to refine theoretical estimates of predictability. We again note that the benefits of as much as 90% of ultimate-duration weather predictions may be available for as little as 10% of the costs of the 'ultimate' system which we've discussed in the foregoing, so (exponentially) demanding are the technical requirements for forecasts of the longest-possible duration.

Climate Change Benefits. Fundamental advances in the understanding of atmospheric physics and weather prediction could also be expected to provide the calibration now lacking for general circulation models of the Earth's response to altered radiative forcing. Those models are now used to predict the impact of anthropogenic effects on time-scales of decades to centuries, despite the fact that current models do not necessarily converge to the long-term circulation patterns actually observed on the Earth when run to very long times. That is, a better understanding of atmospheric physical processes, of real-world mass, momentum and energy flows, and of cloud formation and evolution vis-à-vis both atmospheric science and weather forecasting could also lead to a real scientific basis for long-term climate impact forecasts, thereby permitting discarding of the present-day *ad hoc* approximations.

We emphasize the unique utility of continuously and densely measuring *in situ* the vertical profiles of tropospheric temperature simultaneously with ground- and ocean-surface temperature measurements made by the associated surface-based set of boundary-condition monitoring stations, which we discussed in a little detail three years ago and which we now contemplate linking to the main Weather-Prediction Engine via the same satellite telephony means as is proposed for use by the micro-airship constellation. Such measurements can be expected to swiftly resolve the present-day discrepancy between long-term temperature trends in the lower troposphere versus those on the ground. We also invite particular attention to the peculiarly great ability of the continuously-operating constellation of atmospheric monitoring stations to diagnose cloud physics in a statistically compelling manner, e.g., by space-time-dense *in situ* measurements of the dependence of water droplet size and concentration on position and state variables, throughout large storm systems; most all currently controversial issues, e.g., long-range advective transport of droplets in tropical storm systems, should be objectively resolved by the first few months of reasonably full-scale constellation operation.

The knowledge gains deriving from such advances in space-time-dense, *in situ* measurement could be worth most of one trillion dollars/year of world-wide avoided costs, if such improved understanding led to substantially more cost-efficient means of minimizing anthropogenic aspects of global change than those now contemplated.

A Representative Programmatic Path. We now sketch what we consider to be a representative path for a program that culminates in the system which we've just outlined. Obviously, many alternate routes to the same goal are feasible, and others may be much preferred, for reasons not obvious to us now.

'Breadboard' Prototyping. A proof-of-principle could be attained quickly by attaching to a standard superpressure balloon of a few meters' diameter – i.e., a few kg of payload-mass

—a suitably integrated control processor, temperature, pressure, humidity and insolation sensors, GPS receiver, battery-pack, photovoltaic array and satellite telephone. If inflated for a pressure-altitude of 10 km, this proof-of-principle system will never risk loss via CFIT. Due to the 0.3 bar pressure-altitude selected, it may be expected to avoid an ice-out fate, and thus can be employed to provide arbitrarily finely time-spaced meteorological data for a time-interval of at least a year. Suitably instrumented and programmed, it will even report on its mode-of-demise, whenever this may occur.

'Brassboard' Prototyping While a small number — perhaps a dozen — of these 'breadboard' prototypes are being so exercised and incidentally are returning quite novel scientific datastreams, a chip-set suitable for a 'brassboard' prototype can be assembled, tested and produced in limited quantity. This chip-set would feature a much higher level of integration, with single chips for power management, for full 8-channel GPS reception, for system control and data-management and for satellite telephony. The solar photovoltaic array, GPS antennae and telephony antenna all would be lightweighted and integrated into the balloon structure, and the sensor-module would be ruggedized and miniaturized. Roughly one hundred of these brassboard systems might be deployed for operational evaluation, at various pressure-altitudes between 5 and 15 km so as to endow them with reasonable *a priori* operational lifetimes.

'Steelboard' Prototyping While brassboard field evaluation was taking place — and the central data-collection facility was being more extensively exercised — the main development of the 'steelboard'-level prototyping components would be conducted. These components would be comprised of a single chip implementation of all powered silicon functions, antennae implemented as metal traces and the photovoltaic array as a semiconducting film-structure on the wall of a lightweighted superpressure balloon outfitted with electrically-controlled buoyancy control. The sensor module of the steelboard prototypes would include a microminiaturized water droplet size-and-density spectrometer, and such other environmental sensors as the program sponsor would consider appropriate. At least one thousand of these steelboard prototypes likely would be deployed for operational evaluation, which naturally would even more vigorously exercise the data-collecting facility. [If as many as 10,000 were to be deployed, the resulting data-set quality would support computational weather forecasting to attain as high quality over the entire Earth as that presently available anywhere, as we noted above.]

The steelboard prototyping components would go into mass production as the field evaluation interval concluded, during which time high-rate serial production facilities would be implemented and operated at low initial rates.

Computational Program. Concurrent with the three stages of prototyping of the atmospheric monitoring balloon-stations, one or more pre-existing weather prediction codes would be modified to accept actual atmospheric data in the formats and at the rates being made available at the various prototyping levels. These codes then would be ready to perform on commercially-available computing systems — commercial variants of the ASCI or 'Blue Gene' systems — of appropriate high aggregate processing-rate, as the full-scale system ramped into existence. [As previously noted, as few as 10⁴ balloons deployed — possibly even within the 'steelboard' prototyping level — should support the generation of world-wide weather predictions of the 3-5 day duration and reasonable fidelity now available for only the most-advantaged locations. Computing systems of the corresponding scale are already rather widely available from multiple commercial sources.]

Program Time And Dollar Scales. We expect that roughly one year should suffice for each of the three prototyping phases, so that the full-scale long-range, high-fidelity weather-prediction system could be in operation roughly a half-decade after the program commenced. The cost of the program would rise roughly exponentially in time from a base of ~\$5 M/year to ~\$100 M/year over this half-decade interval.

Risk Management. It is important to note that each of these prototyping and evaluation phases of the proposed program has clearly defined objectives, objective decision criteria for continuation and graceful exit strategies. Sponsor technical and programmatic risk-profiles are stringently minimized in each phase, and the sponsor's forward commitments are for at most a single year's duration. We anticipate that this set of features will be especially congenial to the program-sponsoring community in the U.S.

International Cooperation. We have emphasized from the commencement of our consideration of the prospects for long-range weather-forecasting, e.g., in our presentation to this Symposium three years ago, the great desirability of the fullest-possible international cooperation in the design, development and operation of the type of system that we propose.

We continue to welcome comments and suggestions for improvement in all respects, and we look forward to active collaborations with colleagues from all nations, as this program gets underway.

Long-term weather is intrinsically a truly global phenomenon, and long-term weather-prediction, by the nature of physical law as it's expressed in terrestrial meteorology, *necessarily* develops forecasts of world-wide validity, just as it *demand*s near-real-time knowledge of the state of the atmosphere from all over our planet. For instance, the day-to-day meteorological phenomena in the western U.S. become Europe's weather, less than a week later. Thus, there are few present-day areas of technical endeavor in which international collaboration is more necessary – or in which the large-and-obvious benefits may be more naturally, uniformly and immediately shared among all mankind.


Conclusions. We have presented a top-level discussion of a proposed system for generating high-fidelity, high geographical resolution weather forecasts – necessarily, of global extent – over future time-durations of 2-3 weeks. We believe that it's feasible to realize such a system in a risk-minimized manner on a half-decade time-scale, at a total program cost of perhaps a third-billion dollars. Continuing operations thereafter, for a full-scale system, would involve comparable annual expenditures. Order-of-magnitude smaller continuing expenditures – of the order of \$30 M/year – could 'buy' the real-time geophysical data to support global weather forecasts of 80-90% of the future-time duration of the full-scale system. The benefits associated with such expenditures, as we've pointed out previously, would be several orders-of-magnitude larger: as large as \$100 B/year to the U.S. alone, and 3-4 times that level, world-wide.

Independently, the data that can be gathered uniquely by a distributed, dense collection of airborne and Earth-surface sensors will quickly put global change research on a far more sound scientific basis, as well as resolve some of the outstanding issues currently in play between the several schools-of-thought as to the anthropogenic component of global change.

Thus, there appears to be excellent scientific and societal justifications for near-term commencement of the proof-of-principle phase of the proposed program, naturally with maximal international cooperation.

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

FIGURE 1: Constellation Parameters



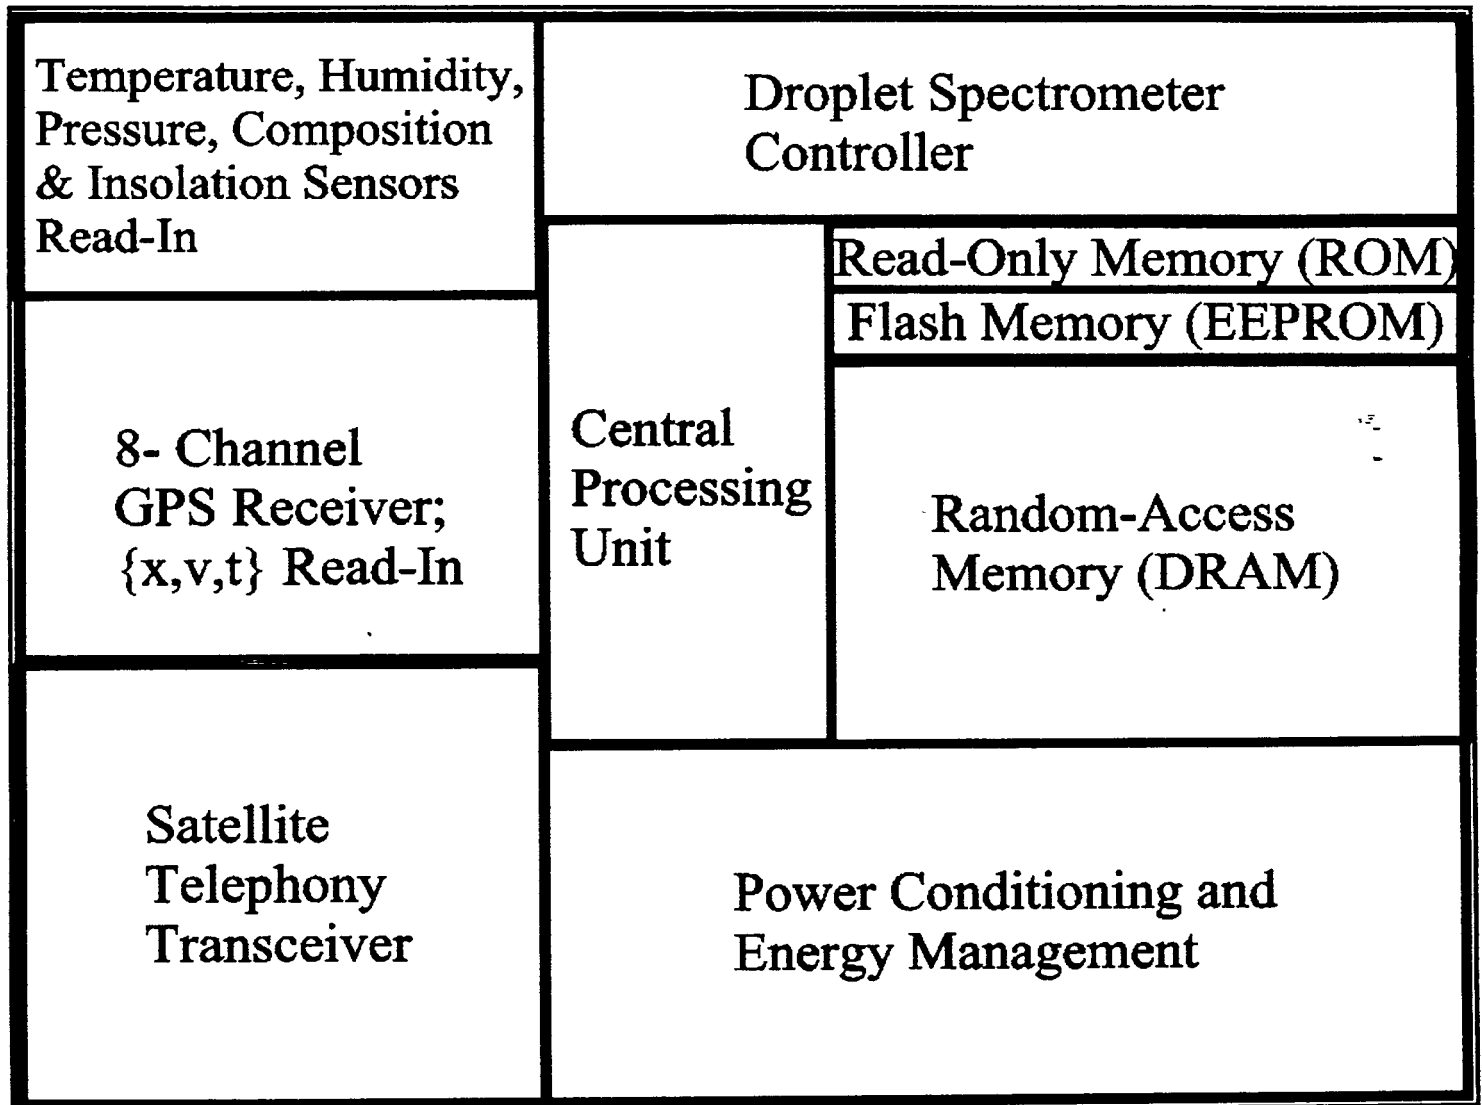


Figure 2. Schematic Layout of Micro-Airship's Single-Chip 'Silicon Payload'




FIGURE 3.