Kent Lindquist White Paper

Number 2001-002

Towards a Successful Anchorage Strong-Motion Network

Kent Lindquist

September 28, 2001

Alaska Earthquake Information Center
Geophysical Institute, UAF

Hello,

We have a complex task ahead, currently involving tangled technical and political issues, in order to make the Anchorage strong-motion network a success. Let me start putting some of my developing answers in print. This paper is an outgrowth of trying to answer a short email of Sept. 26, 2001 from Roger Smith. **If you want the executive summary, read the bold print.**

Roger Smith wrote:

>

> **<u>I think Jeff has a good point here which requires an answer. How good is</u>**
> **<u>Tom's analysis?</u>**

Everything I've seen in print so far misses the main point. DSL vs. IP vs. Serial vs. k22orb vs. k22ew are all possible parts of the first step, yes, namely getting the data into a central collection point. The key issue here is how to achieve robustness. We are limited in that by what the sensor and datalogger can do. Actually the K2 does quite well: it can buffer continuous data for up to 15 minutes, and it can simultaneously collect and store event triggers (segmented data) for days to weeks.[1] That 15-minute continuous-data buffer may be expandable by factors of two, but not orders of magnitude. At best it could be increased to perhaps two hours by sacrificing segmented-event disk space to enlarge the real-time buffer in the K2.

Integrated handling of segmented and continuous data. Ensuring robust data recovery during and after a large earthquake requires a two-pronged approach. First, of course, we need telemetry links between each station and the central collection point that will survive strong shaking. Second, we need integrated handling of the segmented and continuous data from the K2. Any acquisition strategy must either guarantee that telemetry never goes down for more than 15 minutes, or rely on the triggered-data storage capability of the K2 to back up the continuous acquisition. According to a report from the Washington State Military Department's State Emergency Operations Center[2], many parts of the communications infrastructure were overloaded or shut down for up to 24 hours after the Feb. 28, 2001 Magnitude 6.8 Nisqually Earthquake. **This is the key point: we need *integrated handling of continuous and segmented data* from the K2's, in order to use the triggered-data storage ability of the K2's to back up the 15-minute buffering of the datalogger's continuous acquisition mode.**

It is true that k22orb is slightly better matched to the infrastructure we have at AEIC, since it communicates directly with an orbserver. However, because I have written the dataflow connec-

1. In Golden last August, I started to present these technical counterarguments to Alex's system diagrams. I never got a chance to finish, since Harley immediately made absolutely clear that there was no guarantee "we" were buying any more K2's. Harley explained that the ANSS instrument will quite possibly be a completely different instrument, maybe even a newly designed one. The problem with that argument is that any sensor/datalogger combination in the near future will be based on the same combinations of disk-drive and flash-memory solutions and will thus probably have similar data-buffering timescales. Thus, even though I am discussing details of the K2's, I believe the issues themselves transcend the particular choice of sensor and datalogger. Finally, regardless of what the ANSS instrumentation committees decide on strong-motion instruments, what we have now in Anchorage are K2s.
2. "State Incident # 01-0438, Nisqually Earthquake Incident Summary," not for public or press release

tions between earthworm shared-memory rings and orbservers, this is a minor point. Without too much difficulty, we could acquire continuous data with k22ew and bring it from an Earthworm shared-memory ring into an orbserver, then use orb communications between Anchorage and Fairbanks for the continuous data. Or we could bring it into an Earthworm shared-memory ring and use Earthworm import/export software to transmit the data to Fairbanks. Or we could bring it in via k22orb and transfer to an Earthworm shared-memory ring for further processing. There are a number of possible combinations. Whether the data land first in an Earthworm shared-memory ring or an Antelope orbserver is a minor issue. k22orb handles segmented data and continuous data seamlessly, automatically recovering segmented data along with as much continuous data as possible as soon as the telemetry comes back up. k22ew recovers only the continuous data. No solution exists in the USGS plan I got in Golden for using the segmented data to back up the continuous-data acquisition. **Thus, as it exists now, the current system proposal will quite likely fail in a large earthquake.** Less minor also are the ways the k22orb/k22ew choice impacts other system issues, which I will describe in a minute.

The last-mile problem of telemetry is inherently local, not global. Staying focussed on critiquing Tom's analysis, however, the other question to discuss is the type of telemetry connections to the instrument. Many of the solutions I have heard so far suggest choosing a single solution ("DSL" vs. continuously used dialup phone lines etc.) to be applied uniformly to all the Anchorage sites. Telemetry to each site is a last-mile problem. The very fact that the choice of technology is not clear cut shows that the best approach should be to delegate the whole problem to a couple competent engineers, who set up communications between each site and the central concentrator on a case-by-case basis. It is a mistake to try to solve this local problem globally. Of course if the engineers, being competent, find they can use one particular technology at several sites, I'm sure they will do so in the interests of simplicity, robustness, and cost-savings. Certainly it is a mistake to say that a panacea solution for Anchorage will support efforts elsewhere. Some sites may be near good public internet routers, some not. Some will be close enough and some too far from phone switches for DSL. Some will have good radio shots where others won't. Even if a single solution is in principle available at all sites, the reliability may vary. These are things for the engineers to decide at each site. Another advantage of dealing with this local last-mile problem on a case-by-case basis is that a diversity of telemetry strategies may provide more chance that at least some data come in continuously during a large earthquake. There are more issues here that impact the choice of technology, for example all the decisions about whether the datalogger side or the concentrator side initiates connections and re-initiates lost connections; also whether phone company contracts really do prevent continuous occupation of a dial-up line, but this will suffice for the moment.

**<u>> Is it possible to answer that question "cleanly" without
> reference to a political agenda?</u>**

It is only possible to answer that question cleanly if one ignores the full picture. **In order to succeed, we need to straighten out the technical management issues, which are in part political.**

The path from the set of system requirements/constraints to the set of design decisions cannot be done as a collection of one-to-one pairings, requirement 1 implying decision 1; requirement 2 implying decision 2 etc. We are entering the world of *systems engineering*,

which is a discipline in itself. Success requires a *simultaneous, joint inversion* of all requirements to produce a workable whole. Let me demonstrate this. Going back to the initial example of acquisition, Roger H. was right several days ago [Sept. 20 email] in pointing out that "You don't just take the telemetry part and divorce it from the software that acquires the data." Similarly, I'd add that you don't just take the continuous data acquisition decisions and divorce them from the segmented-data acquisition issues. Nor do you divorce the real-time data handling from the offline handling since, as pointed out to me by people in Golden last August, this system has to be able to respond to data submitted after the fact by other people in industry-standard formats. The same calculations will have to be applied to the near-real-time and the historic data files. Also, not to be confused with the buffering support in the datalogger, you do not divorce the selection of packet-format and communication protocols (e.g. Earthworm shared-memory ring vs. Antelope orbserver for data within a few seconds; Earthworm wave-server vs. Antelope orbserver vs. Antelope database vs. Oracle database for data within a few minutes or hours or days) from the buffering strategies at the central collection node. Now is where the choice of intake method for continuous data (k22orb vs. k22ew), described as minor when seen in isolation above, suddenly becomes important. It is also not possible to divorce the choice of buffering strategies at the central collection node in Anchorage from the communication strategy between Anchorage and Fairbanks. The Fairbanks Antelope-based infrastructure integrates short-term buffering into the packet communications itself, via the lossless orbserver communications protocols. The USGS Earthworm-based infrastructure has a small amount of buffering built in to the import/export protocol, however there is no guaranteed losslessness in the protocol. Rather, short-term buffering is provided by the wave-server and its associated client/server protocol. One detail to consider here is that though the wave-server buffers the data, there is no provision in the Earthworm/USGS infrastructure to re-incorporate that buffered data back into the near-real-time packet stream. Yes, the data are available, but via an entirely different protocol. Healing this disconnect is probably possible, however it requires design and implementation effort. Thus, the choice of packet-intake mechanism (which affects k22orb vs. k22ew choice) cannot be separated from the decisions about buffering strategy to back up the Anchorage<->Fairbanks data connection against short-term interruption. That these things are not thought through or designed yet in the USGS/Earthworm plan is clear from an exchange I had with Alex Bittenbinder after our August meeting in Golden.

>[KENT] 1) How do you intend to send the triggered waveform data to Fairbanks?

[ALEX] We have 'sendfile' and 'getfile'. This is a fairly secure means of moving a directory of files from one machine to another. Fairbanks could thus get whatever files come off the K2s.

A second option might be putting traces into the DBMS. Fairbanks could then recover event trace data in SAC via web pages at the EOC.

....except that the USGS/Earthworm plan does not yet have a strategy for getting triggered files from K2's that are operating under continuous mode, according to the diagrams I got in Golden. Nor does it have a strategy for interleaving triggered data from the triggered-only K2's with continuous data. Nor does this proposed 'sendfile' and 'getfile' mechanism mesh at all with our own data intake in Fairbanks. Recovering SAC files would be possible for us. However, we would again have to set up a special-case mechanism for this intake since it is completely separate from the main mechanism of near-real-time transfer. As for the database, there's another disconnect, since the USGS intends to use Oracle[1] and we use Datascope. Those details illustrate some more tangles. What is important to recognize: **The fact that Alex is coming up with these several options as of Aug. 30, and at that only after my prompting, about fundamental, critical design issues indicates that the USGS does not have a complete solution to these problems and will have an extensive amount of design work left even if we agree to do things entirely their way.** They are rapidly designing as they go. Continuing with my email exchange with Alex:

> >[KENT] 2) What happens to the continuous data if telemetry goes down? Are the
> >continuous data lost? Does it depend on the length of the telemetry
> >downtime?
>
> [ALEX] If the k2 -> acquisition link goes down: The K2s have some buffering time, and k2ew will try to recover missing data.
> If the Anch -> Fairbank link goes down, (1) export has a configurable buffer for 'short' breaks. (2) The WaveServer would provide buffering up to months.
> There are WaveServer clients which continuously requests contiguous time slices, or manual ones which recover selected time periods.

This exchange mainly acts to support my previous assertions about weaknesses in the USGS plan as it stands. Also, though, Alex's comments about the 'continuous request of contiguous time slices' reveals (at least to me) how clearly the buffering backup strategy for short-term Anchorage->Fairbanks telemetry outages is tied also to the details of longer-term archiving decisions[2]--file formats, access methods, database types and API's, etc.

I will spare the readers at this point from further elaboration. However, *this is just an introduction.* One cannot divorce the waveform-data transmission issues from the parametric-data exchange issues, since the latter also use parallel channels of communication (orbserver and/or

---

1. Oracle costs thousands of dollars, probably many thousands for the application we are considering. Either we or the USGS or both will have to buy it. I brought this up last August in Golden. The issue was dismissed by Harley on grounds that through ANSS, the USGS will be working out a blanket license for everybody to use Oracle. I would like to point out that this arrangement does not exist yet. More importantly, Oracle requires extensive maintenance effort. We would probably have to hire another person as an Oracle DBA if we maintained it. Harley and Alex's answer to that is that they would maintain the Oracle database out of Golden. Without arguing now the questionable merits of that centralized arrangement for Alaska, there is a further problem: very few seismic analysis applications exist that run off the Oracle platform. Especially, there are none right now that would support AEIC operations. By contrast, the Datascope database has extensive software available on which we base our lab productivity. The point here is not that we need to use Datascope in Anchorage, nor that we absolutely need to do things the AEIC way! The point is that there is a wealth of unaddressed system-engineering detail that goes into these decisions. They look simple on the surface only if one ignores all the critical related issues.
2. Incidentally the archiving strategy Alex proposes is five years behind the state-of-the-art. I know this because that strategy and the supporting software were written at the Geophysical Institute in 1996; rejected out of hand by the USGS; thrown out by us in late 1997 when better technology came along (orb2db from Colorado); and copied in 1999 by the USGS when they realized, belatedly, that they needed an archiving strategy.

import/export) and also need buffering to survive short-term communications loss. The USGS plan shown to me in Golden has a clear inconsistency, admitted to me by the Earthworm team upon questioning, about storing and retrieving parametric data during pre-shakemap signal processing in the shared-memory rings vs. internal to the executables vs. in the Oracle database. Also, there are inconsistencies between the buffering time-constants, decisions on how to ensure losslessness of messages, and access protocols between the USGS and AEIC infrastructures. These inconsistencies need to be healed because the Anchorage shake-map implementation involves a lot more than transmitting K2 data packets from Anchorage to Fairbanks. Fairbanks weak-motion data has to go back to Anchorage, in near-real-time, plus with buffering and recovery strategies for short-term communications loss. AEIC and ATWC earthquake detections have to go to the Anchorage installation to trigger shakemap computation. All these larger issues are in addition to the minor irritations of packet-format and message-format conversions. Finally, importantly, Fairbanks has to provide a backup shakemap product equivalent to the one generated in Anchorage. Even if we put computers with a USGS infrastructure in AEIC itself for this backup task, we still have to resolve the above inconsistencies. The k22orb vs. k22ew choice in Anchorage cannot be made in isolation. It is not as simple as saying that Antelope and Earthworm can exchange packets. All of these issues must be considered simultaneously to do a joint-inversion for a complete technical plan.

The joint inversion, i.e. systems engineering for a working technical blueprint, cannot be done in the Markov approximation: it cannot be done while ignoring past work and future directions. This inversion for a working system blueprint is *hysteretic with a multi-year time constant.* Stated alternately, these details are complex and represent an *infrastructure*, built with years of investment, that makes current performance and future progress possible. We at the GI have invested over a decade in building a seismic monitoring infrastructure. Only a fraction of this has been funded by the USGS. I myself have committed six years of my career to infrastructural development of the seismic monitoring system at the Geophysical Institute. The point here is not to imply that we should be guided by our investment (of course, we should be guided only by our potential to gain). Rather, it is to show that the infrastructure to handle all the myriad details of modern, advanced seismic monitoring gets quite complex. Such an infrastructure must be nurtured carefully and engineered with well thought-out systems-level decisions. The USGS software contributions have definite merits and solid contributions, many of which we use at the GI. However, as a package the USGS infrastructure alone comes nowhere near to meeting our needs as a regional Alaskan monitoring network. Don't learn this the hard way. We cannot throw out our infrastructure without failing as an organization, i.e. failing to provide the legislatively-mandated monitoring services and science support that we do now.

The first temptation is to suggest we are suffering from "not invented here." In fact, for years we have taken the approach of gluing together components from as many sources as possible to meet our full needs. This approach is the antithesis of "not invented here." The Iceworm system we are running now began as a 50/50 mixing of USGS Earthworm software and IRIS/JSPC Datascope software, the latter of which evolved into Antelope. We continue to use many pieces of USGS software now, including the main Earthworm associator for automatic locations of events. What is happening is that the USGS architecture and the AEIC architecture are evolving in distinct and not always compatible directions, exactly at the time when the interconnections between the two infrastructures are getting more complex. The true questions are not between a USGS software module vs. an Antelope software module, or a USGS packet protocol vs. an AEIC packet protocol or format. The true issues: what is the appropriate path for the joint evolution of these infrastructures? In order to continue to deliver current products and expand to deliver more advanced products, we need a carefully engineered approach to future developments. **The operative question now becomes political: Alongside the USGS infrastructure, does the GI seismic monitoring infrastructure (and that of AVO, and that of ATWC) have a right to coexist?** This question must be answered, and that answer must be actively supported at all levels, in order for any technical collaboration to succeed. I believe the unequivocal answer to this question, consistent with ANSS literature supporting regional autonomy and regional/ national partnership, must be 'yes'. **Key to this coexistence is developing a commitment to open-systems architecture** in the U.S or at least in the Alaskan seismic community. Output products, communications protocols and formats, and perhaps algorithms should be specified. The rest should not. I believe such a commitment to open-systems architecture is necessary to the success of ANSS, lest the system become archaic before it is installed, lest the system fail to benefit from distributed academic expertise, lest the system fail to adjust for regionally varying needs, lest the system fail to give itself the flexibility for ongoing growth.

As of this moment the Anchorage Strong-motion network development is still floundering. This is fixable. We have established a cohesion of purpose internal to the GI, a joint agreement that our seismic monitoring efforts will continue, that our infrastructure has a right to coexist, and that we are positioned to play a major role if not the lead role in inter-agency seismic monitoring within the state. **Internal cohesion is necessary but not enough. We now need to develop a clear, articulated vision of how seismic monitoring is conducted in the entirety of the state of Alaska.** This master plan does not preclude the involvement of other agencies. Rather, it relies on it. Nevertheless, we will fail without such a master plan. **This introduces the second political issue: we need to develop this master plan collectively as a set of Alaskan monitoring agencies jointly responsible for monitoring the state.** We need to decide who monitors what parts in what ways, who is responsible for which output products, who archives what etc. We also need to pool resources to support these efforts. We are in effect pooling resources now, albeit in an inadequately coordinated way, in the Anchorage strong-motion implementation. **As a first step, I suggest that the GI internally develop a strawman master plan for how seismic monitoring should happen in Alaska, then negotiate amongst other Alaskan agencies to come to consensus** on how to meet all of our needs. The technical complexity of engineering a monitoring solution for Anchorage strong-motion is showing us a new reality: we are no longer a set of completely independent monitoring agencies in Alaska, with infrastructures that can be considered as modular blocks with a few connections between them. Our missions are growing together and the connections amongst our infrastructures are rapidly becoming more complex. **Thus, we need to start think of Alaskan seismic monitoring as a whole, and engineering consistently across the state with that in mind.**

Developing such a consensus political and technical vision of how to conduct monitoring in Alaska is the first component of something we, as a community, need to develop: better technical management. The Anchorage strong-motion implementation is currently floundering because the complexity of the task exceeds the performance of our community's technical management strategies. The first step is to develop an internal GI vision of state of Alaska monitoring, then share it with other Alaskan monitoring agencies and negotiate our way to a consensus. **The second step, the second technical management issue, is to develop better delegation**. Just as the engineers should not be telling the managers what the priorities are, the managers should not be telling the engineers what telemetry links and software to use. Managers should not be burdened with deciding amongst DSL vs. phone lines. Managers should not be burdened with deciding amongst k22orb and k22ew. Managers should certainly not be burdened with doing the actual configuration work for any of these systems or for prototype systems tests. **Rather, all of these issues should be delegated to the engineers, freeing up the managers to lay out organizational structure and to decide how to pool resources to pay for programmatic successes.** The engineering looks simple on the surface. It is not. The full complexity needs to be respected, and delegated, to people or teams (not committees) of people who can see the full technical picture. Delegation means not only transferring accountability [as has been done so far], but also transferring authority to make decisions and transferring ownership of the task. Management needs to actively get out of the way, i.e. stay out of the details.

I hope these comments are taken in a constructive light. I am not targeting any one person. Rather, from my experience on ANSS committees and subcommittees, I believe that not just Alaska but the U.S. seismic community as a whole is suffering from these problems. **We are trying to run what will ultimately be a multi-million-dollar program by 'winging it' on the technical management. We need to correct this or we will all fail.** The fact that we do not yet have the millions of dollars, and are instead scrambling for resources donated from other budgets, makes us more sensitive to this issue, not less.

Regarding the establishment of a joint monitoring system for Alaska, through ANSS the USGS implicitly proposes to do this all for us. However, they do not currently have the resources, requiring instead a grassroots approach, building on what already exists and developing partnerships. This, actually, is extremely healthy, since it is in keeping with the philosophy of supporting regional autonomy and collaboration, with which ANSS has been presented at high levels. Succeeding at developing a true partnership between Alaska as a region and the USGS efforts is certainly possible, and will benefit us all. As my respected colleague Dan McNamara has said, "There's enough room in the world for everyone." Or, as restated by Jeff Freymueller, "There's more than enough work for everyone." **This development is happening in an environment where the various Alaskan agencies are working hard to assemble an effective collaboration. In order for this to succeed, we need to make sure that there is extensive communication and coordination amongst all Alaskan parties and the USGS, lest uncoordinated activities have the effect of driving wedges amongst the players precisely as we are working to draw together.**

While the USGS and other parties at the national level are still pursuing funding and support for ANSS, and while a number of national committees are still working to put together guidelines for ANSS technology (deliberately recusing themselves, according to their reports, from even considering actual details at this point), **we in Alaska are already at the forefront, implementing a pace-setting, predecessor ANSS system** in the Anchorage strong-motion network. **We have a stellar opportunity to create a joint achievement between the USGS and Alaska, showing an example of how autonomous regions and USGS Nationwide can succeed in implementing ANSS.** We need only to solve these several technical, political, and organizational issues in front of us, and we will all succeed in our respective missions.