System Slack in Cyberinfrastructure Development: Mind the Gaps

Nicholas M. Weber and Karen S. Baker
University of Illinois, Champaign- Urbana
Graduate School of Library
and Information Science
nmweber@illinois.edu; ksbaker2@illinois.edu

ABSTRACT

In the following paper we draw on Leigh Star's notion of infrastructural orphans [7], those awkwardly aligned with the technology of an emerging infrastructure, to offer a cyberinfrastructure (CI) developmental principle referred to as 'system slack'-- the capability designed into an information facility to vary in it's tolerance or adherence to designated standards. We argue that if CI's are to be developed with the capability of *simultaneously* coordinating large-scale data collaborations and local heterogeneous practices then developers and designers must keep in mind that there always will be some elements within the configuration of a CI system that are *momentarily* awkwardly aligned. Gaps in connectivity will result in under-performing, or infrequently occurring CI's.

Author Keywords
Cyberinfrastructure development, design principles, systems development, standards.

SOCIOTECHNICAL ARRANGEMENTS

Although a single definition of CI is elusive given the complexity and breadth of the topic, we rely on Edwards et al's cogent summary; "Cyberinfrastructure is the set of organizational practices, technical infrastructure and social norms that collectively provide for the smooth operation of scientific work at a distance." [3; p.6] However, we note that in other literature CI is frequently presented as a single entity or single "solution" to networked connectivity issues existing between projects various domains of scientific knowledge [1]. In practice CI is a distributed set of arrangements composed of a suite of components tied to each other via a web of connections. Taking a multi-component view, we consider two or more information facilities connected by digital technologies, associated participants and their agreed upon standards to be the basic elements of an example *CI system*.

Where infrastructures are said to be built upon an installed base [6], cyberinfrastructures are built upon installed *bases.* A CI is necessarily positioned between or among two infrastructures- if not more. Referring back to the Edwards et al's summary; it is the practices, infrastructures and social norms *collectively* that makeup a CI. These nested arrangements in a system of systems contribute to a CI's conceptual complexity, making it difficult to balance design requirements in total ('end-to-end') with those more specific to individual components. Non-hierarchical configurations have proven particularly challenging for CI development that attempts to resolve tensions between local and global scales of collaborative work so that an infrastructure can 'occur' [6]. As Anna Gold [4] noted about cyberinfrastructures generally, they "...tend to emphasize global, large-scale data-taking and data-management enterprises, and these neither exhaust nor address the full extent of science data production and use in knowledge-making. Complementing these solutions are others that focus on local knowledge-making and local data centers." Moving beyond a hierarchical arrangement requires that we carefully consider the features of this web of connections, so that we can more capably address the 'full extent of scientific data production and use.'

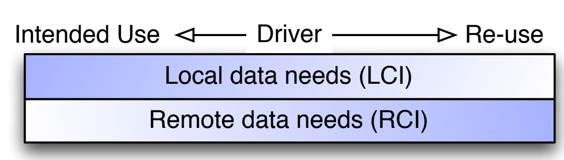

Figure 1. These two paired gradients depict the priorities driving the work of an information facility, where local and remote data needs make inverse demands for system development and capability.

Features of Nodes

Integral to a non-hierarchical view of CI are what we are calling 'information facilities', communicating nodes that anchor a web of internet-based connections. Each facility might be described as an organizational entity or arrangement that provides *support* for the storage, access and/or organization of data in a networked collaboration. At

CSCW'12, February 11–15, 2012, Seattle, Washington, USA.

some levels this support will be focused on the practicalities of project-based scientific practice where coordinating data use among project or team members will be the biggest priority. At other levels, this support will be directed towards the re-purposing or reuse of data, such as model construction, or other collaborative scientific work removed from the original context in which the data were generated [1].

Traditionally distinctions between information facilities have been made with titles like ‘data centers’ or ‘computational facilities’ but these terms hold little explanatory power in the context of CI.. In Table 1 we present some of the distinguishing features that characterize information facilities. The table presents four exemplar nodes from a continuum of facility types.

Features of Nodes (Information Facilities)				
	Lab Based	*Multi-Project Based*	*Community Based*	*Multi-Community Based*
Socitotechnical distance from data origin and focus	Local (micro)	Local (meso)	Remote (meso)	Remote (macro)
Individual, production	Individual	Individual/Production	Production	Production
Focus	Centralized	Centralized	Distributed	Distributed
Data Relation	Problem-based	Multi-Project-based	Network-based	Network-based
Adherence to standard	Targeted	Targeted/Single Community	Single/Multiple Community	Multiple Community

Table 1. Information facilities: Four exemplars from a continuum of types.

Important to the framing of the table is an intertwined pair of drivers that prioritize work: problem-based data use in local cyberinfrastructures (LCI) and network-based data reuse typically associated with remote cyberinfrastructures (RCI) (see Figure 1). The influence of these drivers may be imagined as gradients operating inversely with respect to each other. At project-based facilities, local data needs take priority in the data work at hand, but the need for eventual coordination with remote network-based facility arrangements is an always-present requirement taken into account in long-term plans and designs. The inverse is true for remote network-based facilities. In practice these facilities seem to conceptually mirror one another even though they are functionally quite different.

CYBERINFRASTRUCTURE TUNING

The purpose of a CI’s deployment should dictate its extensibility, i.e. the amount of variation it’s designed to accommodate. Bietz et al [2] describe a frequently overlooked point about CI development- it is often built with a very specific purpose in mind, though it must also, by it’s very purpose in facilitating connectivity, be interoperable across unimagined layers of infrastructure. The greater the requirements for complete fidelity to *community standards*, the more narrow or finely tuned the CI’s focus becomes, and this can potentially isolate local arrangements.

LCI’s typically have a commitment to more targeted standards that have been developed with the aim of adjusting systems to meet local practices and needs. These local infrastructures are often ad-hoc in their construction, yet agile and flexible in accommodating changes in practice since their design approach is tuned to specific data-needs and support of local data practices.

As noted earlier, local practices are often project specific and meant to accommodate a practice of science that is, as Knorr-Cetina [5] suggests, essentially about tinkering and heuristic innovations to ‘make things work’. One measure of CI success may be the extent to which partners are not only aware of orphan data held locally, or stored in database systems, but also are collaboratively active in planning for awkwardly aligned elements. Indeed, the recognition and acceptance of existing data issues, adds validity to local articulation and innovation. Such a process ensures continuing update of an ‘end-to-end’ cyberinfrastructure.

Principle of System Slack

The capability designed into an information facility to vary in it’s tolerance or adherence to designated standards is what we call a *system’s slack*. In developing this design principle we will discuss how standards are used at a 'targeted' level by meso-scale facilities focused on connecting problem-based projects and PI's; these facilities will necessarily have a large amount of slack built into their accommodation of ad-hoc or emergent standards. But we will also consider standards that are necessarily 'bounded' at macro-level facilities committed to enabling large-scale collaborations with more well-defined data sources committed to community wide standards. Our initial description of standards use and designing system slack in these two types of facilities is purposefully generic -- we want to suggest that these are two ends of a spectrum that contribute equally and will require balanced consideration for future CI development. However, we suggest that less is understood about the diversity of LCI approaches needed for accommodating a wide variety of actors and their practices.

For example in an LCI, data 'of a kind' may be incorporated into a developing data system while data 'misfits' or orphans are often put to the side to accumulate and await further attention from data managers. The misfits of an information environment are often misaligned in some way with the standards of 'good-fit' data, but not so incommensurable that they can't be ingested by the system for future integration. These may be stored as stand-alone legacy files online or incorporated in separate, independent systems of the information facility but are not integrated into the cyberinfrastructure itself. Franklin et al [8] have highlighted the value of associated but independent databases.

The misfits (or orphans) and their relative size and importance serve as prompts for developers considering whether to modify an existing system or to design a new system- they represent a need to 'figure things out'. On one hand, designers may choose to modify information system specifications in order to accommodate new types of information. This may be made possible by changing constraints and capabilities relating to technical programming and/or architecture possibilities. If incorporated within a local, targeted system that adheres to a standard delivery, it is implicitly understood in accepting them that the newly incorporated data will be cast into 'standard' formats in order to be made more widely accessible to project partners or disseminated broadly in existing network arrangements. Recasting new types of data may or may not be problematic. The process may make evident new requirements not accommodated by the existing standard in which case more 'figuring out' must be undertaken. Often, when the fit between standards and data is not good, the data may be 'shoehorned' temporarily into standard format anyway. In many cases a poor solution wrangles in the minds of those involved and the discontent is the stimulus for future innovation. These fitting, calibrating, accommodating and figuring out activities are seminal to designing interoperable systems, but at the LCI become exaggerated by the need to accommodate very diverse data types and support ever-evolving heterogeneous practices. In our full paper, we will describe a salient case that illustrates how the failure to account for system slack represents a failure of the CI with respect to the potential success of integrating ecological data.

FINAL THOUGHTS

As Star and Ruhleder have insightfully stated, "An infrastructure occurs when the tensions between the local and global have been resolved" [6, p. 114] and that these relations between local and global will always be momentary as no "home is universal." This tenet is especially true for 'occurrences' of a CI where awkwardly aligned data or the incomplete use of standards may be momentarily outside a given systems capability of interoperation. We have argued above, that for the design of sustainable, ongoing CI's there needs to be an account for the tinkering of scientific practice which will *always* produce some data that are awkwardly aligned with a system's standards- no matter the level of aggregation or the breadth of collaboration that is attempting to be brokered. The ability for a system to account for these variations in standards' use, a system's slack, is not a rule nor even a standard itself per se, but instead a principle for design- one that allows the promise of cyberinfrastructure to be realized through the practices of continuing alignment.

REFERENCES

1. Atkins, D. et al. (2003). " Revolutionizing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-ribbon Advisory Panel on Cyberinfrastructure" NSF, Report of the National Science Foundation Blue-ribbon Advisory Panel on Cyberinfrastructure. Retrieved November, 2011: http://www.nsf.gov/od/oci/reports/atkins.pdf

2. Bietz, M. J., Baumer, E. P. S., & Lee, C. P. (2010). Synergizing in Cyberinfrastructure Development. *Computer Supported Cooperative Work*, *19*(3-4), 245-281

3. Edwards, P. N., Jackson, S. J., Bowker, G. C., and Knobel, C. P. (2007). *Understanding Infrastructure: Dynamics, Tensions, and Design*. Ann Arbor: Deep Blue

4. Gold, A., 2007. Cyberinfrastructure, Data, and Libraries, Part 1, A Primer for Librarians. D-Lib Magazine 13 (9/10).

5. Knorr-Centina, K., 1979. Tinkering Towards Success: Prelude to a Theory of Scientific Practice. Theory and Society 8: 347-376.

6. Star, S.L. and S.Ruhleder, 1996. Steps Toward an Ecology of Infrastructure Design and Access for Large Infrastructure Spaces. Information Systems Research 7(1): 111-134.

7. Star, S.L. (2007, March) 'Orphans of Infrastructure: A New Point of Departure' The Future of Computing: A Vision. Lecture conducted from Oxford Internet Institute, Newham College, Cambridge, UK.

8. Franklin, M. Halevy, A. and Maier, D. (2005) From databases to dataspaces: A new abstraction for information management. *Sigmod Record*, 34(4):27–33, 2005.