Table 1
Mathematical Notations

Symbol	Description
P_θ	Prior probability distribution of events
$P_{\theta,n}$	Prior probability of the number of events
$P_{\theta,t}$, $P_{\theta,d}$, $P_{\theta,l}$, $P_{\theta,m}$	Prior probability of event time, depth, location, and magnitude
e	Set of events; e^i is the ith event
e^i_t, e^i_d, e^i_l, e^i_m	Event time, depth, location, and magnitude
P_ϕ	Probability of arrivals given an event
$P_{\phi,d}$	Detection probability of an event
$P_{\phi,z}$, $P_{\phi,s}$, $P_{\phi,a}$, $P_{\phi,h}$	Probability of arrival azimuth, slowness, amplitude, and phase given event
Λ	Set of arrivals associated to events
Λ^{ijk}	The arrival of event i's jth phase at station k
Λ^{ijk}_t, Λ^{ijk}_z, Λ^{ijk}_s, Λ^{ijk}_a, Λ^{ijk}_h	Arrival time, azimuth, slowness, amplitude, and phase
ζ	A special symbol that signifies the lack of an arrival
P_ω	Probability of false arrivals
$P_{\omega,n}$	Probability of number of false arrivals
$P_{\omega,t}$, $P_{\omega,z}$, $P_{\omega,s}$, $P_{\omega,a}$, $P_{\omega,h}$	Probability of false arrival time, azimuth, slowness, amplitude, and phase
ξ	All false arrivals. ξ^k are the false arrivals at station k, and ξ^{kl} is one of these false arrivals
ξ^{kl}_t, ξ^{kl}_z, ξ^{kl}_s, ξ^{kl}_a, ξ^{kl}_h	False arrival time, azimuth, slowness, amplitude, and phase
P_γ	Probability distribution of coda arrivals
$P_{\gamma,d}$	Probability that a coda arrival is detected
$P_{\gamma,z}$, $P_{\gamma,s}$, $P_{\gamma,a}$	Probability of coda azimuth, slowness, and amplitude given previous arrival
$P_{\gamma,h}$	Probability of coda phase
η^a	Coda arrival generated by arrival a or ζ if no coda
η^a_t, η^a_z, η^a_s, η^a_a	Coda arrival's time, azimuth, slowness, and amplitude
η	Set of all coda arrivals
A	Set of all observed arrivals
Δ_{ik}	Great-circle distance between event i and station k

historical data and present an algorithm that infers the set of seismic events given the observed detections. In simple terms, let X be a random variable ranging over all possible collections of seismic events, with each event defined by time, location, depth, and magnitude. Let Y range over all possible signal observations at all seismic stations. Then $P_\theta(X)$ describes a parameterized generative prior over events, and $P_\phi(Y|X)$ describes how the signal is propagated and measured (including travel time, selective absorption and scattering, noise, artifacts, sensor bias, sensor failures, etc.). Given observations $Y = y$, we are interested in the posterior distribution over events, $P(X|Y = y)$, which is given by Bayes rule:

$$P(X|Y = y) \propto P_\theta(X)P_\phi(Y = y|X).$$

For the CTBT monitoring application, we are interested in the value x^* that maximizes this expression, that is, the most likely explanation for all the sensor readings:

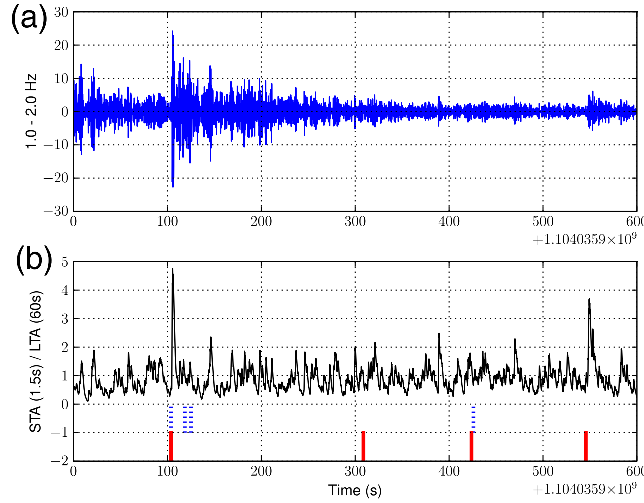

Figure 1. Example of seismic waveform (station ASAR, channel SE), STA/LTA, and arrivals. At the bottom of the lower panel the dotted lines are the automatically detected arrivals, while the solid lines are the analyst marked arrivals. The color version of this figure is available only in the electronic edition.

$$x^* = \arg\max_x P_\theta(X = x)P_\phi(Y = y|X = x).$$

Assuming that an algorithm can be devised to solve this optimization problem, which turns out to be nontrivial, the key determinant of the success of the Bayesian approach is the accuracy of the probability models P_θ and P_ϕ. These models embody, in explicit form, elementary knowledge of seismology and seismometry, as well as the residual uncertainty inherent in that knowledge. Adding to and refining the knowledge embodied in the models reduces residual uncertainty and improves system performance.

Many other researchers have previously applied Bayesian techniques to geophysical problems (see Duijndam, 1988a,b for examples). Most of these applications involve inference over a fixed set of continuous-valued parameters; the event localization problem in particular is addressed by Myers *et al.* (2007). The full monitoring problem involves inference over a combinatorial discrete space (the number of events and the association between events and observations), as well as a continuous parameter space for each event. In this respect, it resembles the data association problems arising in multitarget tracking (Bar-Shalom and Fortmann, 1988).

Our overall project, VISA (vertically integrated seismic analysis), is divided into two stages. The first stage, NET-VISA, is the subject of the current paper. As the name suggests, NET-VISA deals only with network processing and relies upon the IDC's pre-existing signal detection algorithm, which converts the raw waveforms into a sequence of arrivals. Thus, the observations Y correspond to sets of arrivals with measured attributes rather than raw waveform signals. An example of the signal processing is shown in Figure 1. Figure 1a displays a filtered seismic waveform, and Figure 1b shows the typical short-term average to long-term average